SPEDGE

Implementing Cisco Service Provider Next-Generation Edge Network Services

Version 1.0

Student Guide

Text Part Number: 97-3156-01

https://t.me/learningnets



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

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

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

DISCLAIMER WARRANTY: THIS CONTENT IS BEING PROVIDED “AS IS.” CISCO MAKES AND YOU RECEIVE NO WARRANTIES IN CONNECTION WITH THE CONTENT PROVIDED HEREUNDER, EXPRESS, IMPLIED, STATUTORY OR IN ANY OTHER PROVISION OF THIS CONTENT OR COMMUNICATION BETWEEN CISCO AND YOU. CISCO SPECIFICALLY DISCLAIMS ALL IMPLIED WARRANTIES, INCLUDING WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT AND FITNESS FOR A PARTICULAR PURPOSE, OR ARISING FROM A COURSE OF DEALING, USAGE OR TRADE PRACTICE. This learning product may contain early release content, and while Cisco believes it to be accurate, it falls subject to the disclaimer above.

© 2012 Cisco and/or its affiliates. All rights reserved.

https://t.me/learningnets



Students, this letter describes important course evaluation access information!

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

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

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

Sincerely,

Cisco Systems Learning

https://t.me/learningnets

https://t.me/learningnets

Table of Contents

Course Introduction 1

Overview 1
Learner Skills and Knowledge 2
Course Goal and Objectives 3
Course Flow 4
Additional References 5
Cisco Glossary of Terms 5
Your Training Curriculum 6
Your Training Curriculum 7

VPN Technologies 1-1

Overview 1-1
Module Objectives 1-1

Introducing VPNs 1-3

Overview 1-3
Objectives 1-3
VPN Concept 1-4
VPN Models 1-12
Summary 1-30

Introducing MPLS VPNs 1-31

Overview 1-31
Objectives 1-31
MPLS VPN Architecture 1-32
MPLS VPN Routing 1-42
MPLS VPN Forwarding Mechanisms 1-48
Summary 1-54
Module Summary 1-55
Module Self-Check 1-57
Module Self-Check Answer Key 1-60

MPLS Layer 3 VPNs 2-1

Overview 2-1
Module Objectives 2-1

Implementing MPLS Layer 3 VPN Backbones 2-3

Overview 2-3
Objectives 2-3
Virtual Routing and Forwarding 2-4
Example: BGP Route Propagation—Outbound 2-11
Enabling VRF 2-18
Enabling MP-BGP 2-32
Summary 2-45

Connecting Customers Using Simple Routing Protocols 2-47

Overview 2-47
Objectives 2-47
PE-CE Routing 2-48
Summary 2-66

Connecting Customers Using BGP or OSPF 2-67

Overview 2-67
Objectives 2-67
OSPF as the PE-CE Routing Protocol 2-68
BGP as the PE-CE Routing Protocol 2-83
Troubleshooting MPLS VPNs 2-94
Summary 2-101

https://t.me/learningnets

Deploying IPv6 and MPLS 2-103
Overview 2-103
Objectives 2-103
Support for IPv6 in MPLS 2-104
Pros 2-106
Cons 2-106
Pros 2-108
Cons 2-108
Description of Pros and Cons 2-108
Multiprotocol Extensions for BGP4 2-110
Summary 2-123
Module Summary 2-125
References 2-125
Module Self-Check 2-127
Module Self-Check Answer Key 2-129

Complex MPLS Layer 3 VPNs 3-1

Overview 3-1
Module Objectives 3-1
Implementing Complex MPLS Layer 3 VPNs 3-3
Overview 3-3
Objectives 3-3
Overlapping VPNs 3-4
Central Service VPNs 3-12
Managed CE Router Service 3-22
Summary 3-25
Implementing Internet Access and MPLS Layer 3 VPNs 3-27
Overview 3-27
Objectives 3-27
Internet Access Models with MPLS VPNs 3-28
Separate Internet Access and VPN Services 3-37
Internet Access as a Separate VPN 3-43
Summary 3-51
Introducing MPLS Interdomain Solutions 3-53
Overview 3-53
Objectives 3-53
MPLS Interdomain Solutions 3-54
CSC Models 3-63
Inter-AS 3-69
Summary 3-75
Module Summary 3-77
Module Self-Check 3-79
Module Self-Check Answer Key 3-81

Layer 2 VPNs and Ethernet Services 4-1

Overview 4-1
Module Objectives 4-1
Introducing Layer 2 VPNs 4-3
Overview 4-3
Objectives 4-3
Layer 2 VPN Overview 4-4
Summary 4-24

https://t.me/learningnets

Introducing AToM 4-25
Overview 4-25
Objectives 4-25
Introduction to AToM 4-26
AToM Implementation 4-47
Summary 4-55
Implementing VPLS 4-57
Overview 4-57
Objectives 4-57
VPLS Overview 4-58
Implementing VPLS and H-VPLS 4-68
Summary 4-74
Module Summary 4-75
References 4-76
Module Self-Check 4-77
Module Self-Check Answer Key 4-78

https://t.me/learningnets

https://t.me/learningnets

Course Introduction

Overview

The *Implementing Cisco Service Provider Next-Generation Edge Network Services* (SPEDGE) 1.0 course is an instructor-led course presented by Cisco Learning Partners to their end-user customers. This five-day course provides network engineers and technicians with the knowledge and skills necessary to implement and support a service provider network.

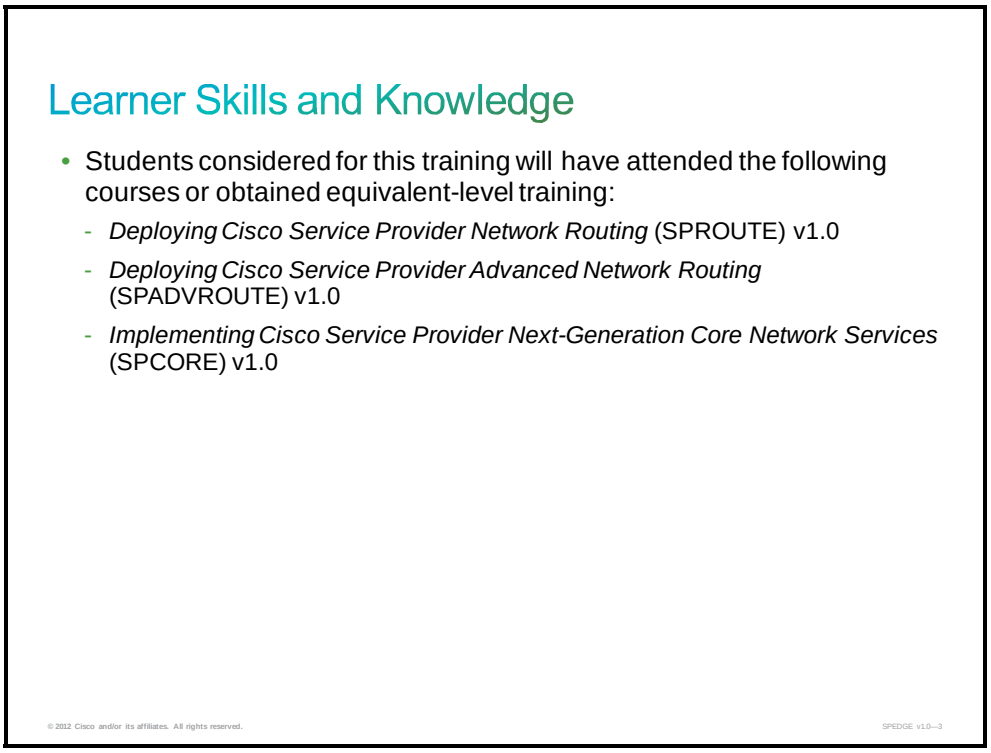
The course is designed to provide service provider professionals with information on the use of service provider VPN solutions. The goal is to train professionals to enable service provider points of presence (POPs) to provide Layer 2 and Layer 3 VPNs. The course reinforces the learning by providing students with hands-on labs to ensure that they thoroughly understand how to implement VPNs within their networks.

The course also includes classroom activities with remote labs that are useful to gain practical skills in deploying Cisco IOS, IOS XE, and IOS XR features to operate and support service provider VPN solutions.

https://t.me/learningnets

Learner Skills and Knowledge

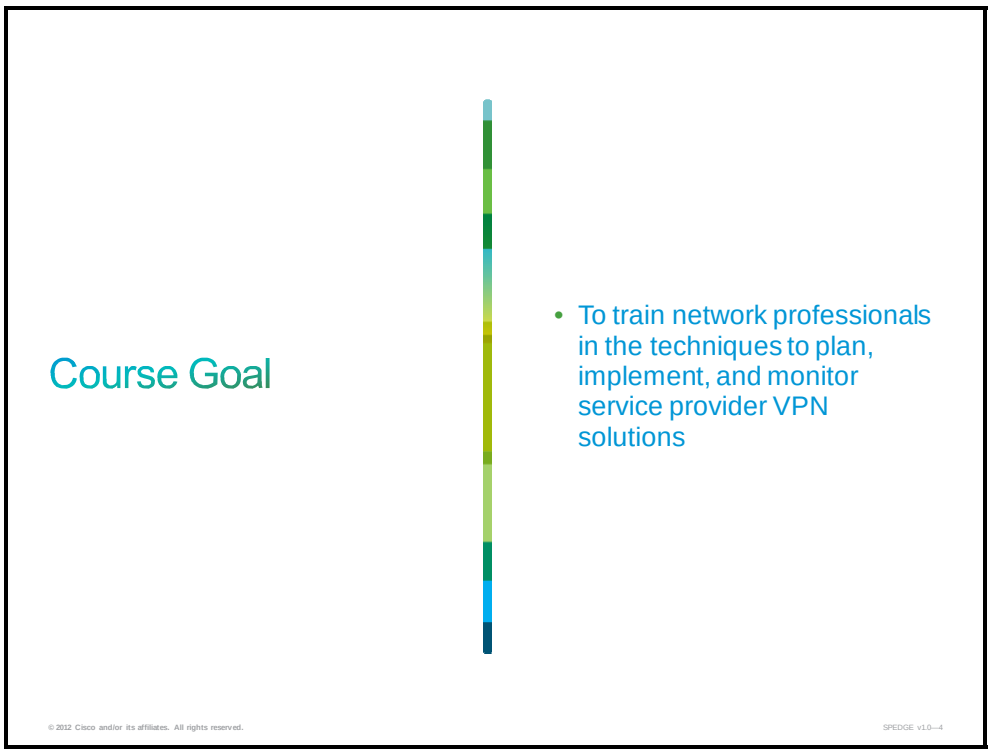
This subtopic lists the skills and knowledge that learners must possess to benefit fully from the course. The subtopic also includes recommended Cisco learning offerings that learners should first complete to benefit fully from this course.



https://t.me/learningnets

Course Goal and Objectives

This topic describes the course goal and objectives.



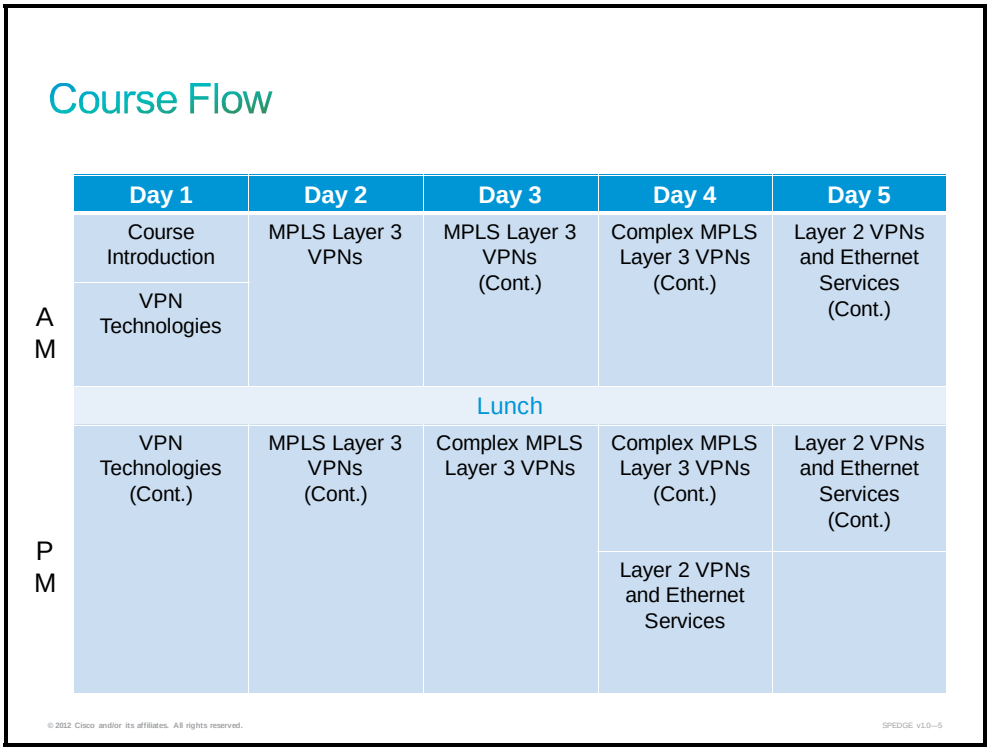
Upon completing this course, you will be able to meet these objectives:

- Introduce the VPN technologies that are used in the service provider environment and the MPLS VPN peer-to-peer architecture
- Describe the implementation steps that are needed to provide MPLS Layer 3 VPN service in the service provider network
- Describe how the MPLS Layer 3 VPN model can be used to implement managed services and Internet access
- Describe Layer 2 VPNs and Ethernet services

https://t.me/learningnets

Course Flow

This topic presents the suggested flow of the course materials.

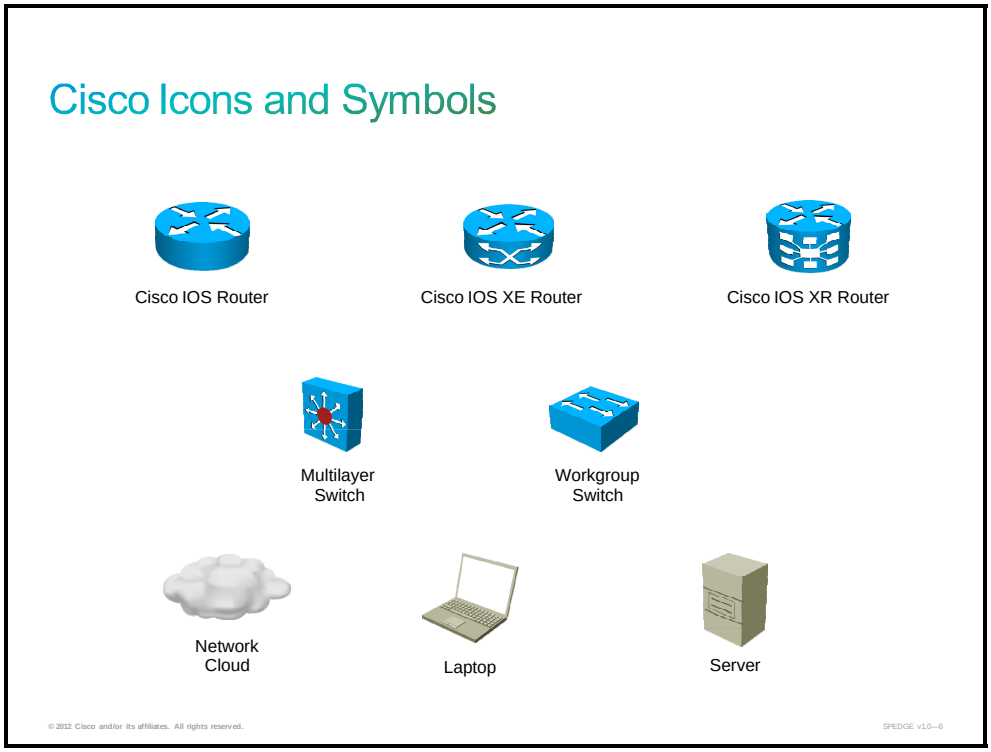


The schedule reflects the recommended structure for this course. This structure allows enough time for the instructor to present the course information and for you to work through the lab activities. The exact timing of the subject materials and labs depends on the pace of your specific class.

https://t.me/learningnets

Additional References

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



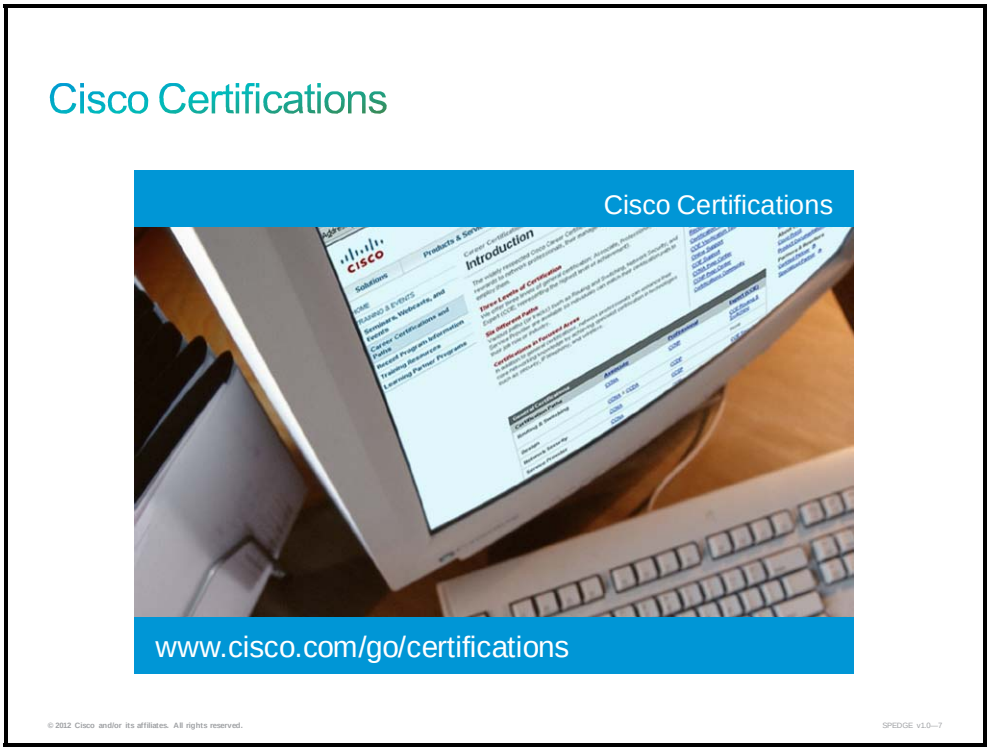
Cisco Glossary of Terms

For additional information on Cisco terminology, refer to the *Cisco Internetworking Terms and Acronyms* glossary of terms at http://docwiki.cisco.com/wiki/Internetworking_Terms_and_Acronyms_%28ITA%29_Guide.

https://t.me/learningnets

Your Training Curriculum

This topic presents the training curriculum for this course.

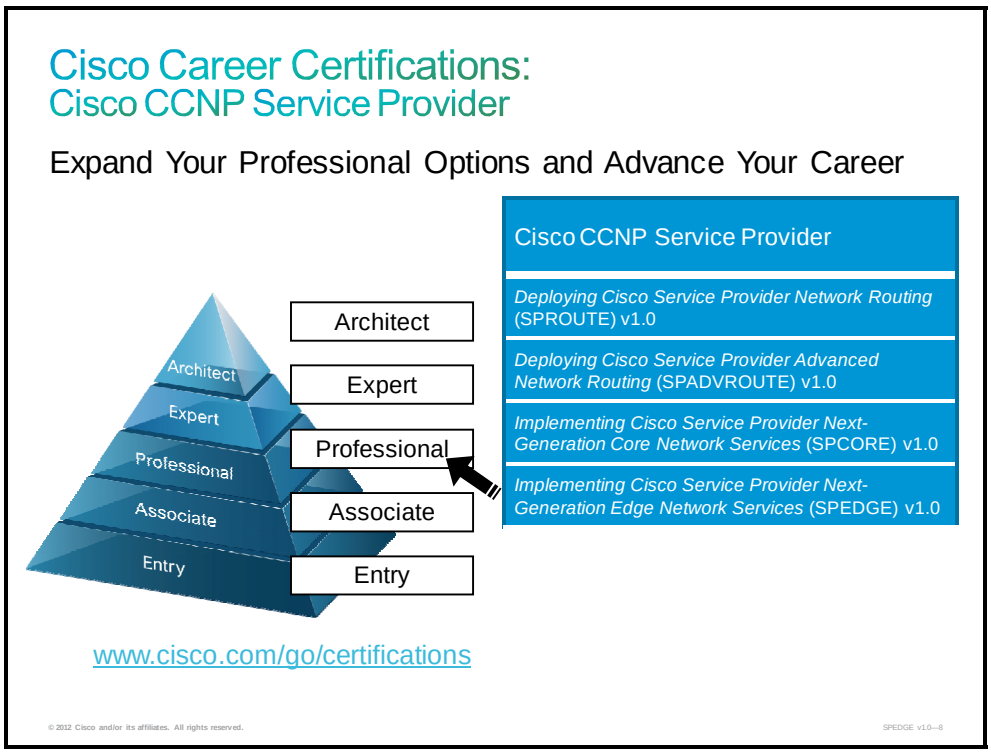


You are encouraged to join the Cisco Certification Community, a discussion forum open to anyone holding a valid Cisco Career Certification (such as Cisco CCIE®, CCNA®, CCDA®, CCNP®, CCDP®, CCIP®, CCVP®, or CCSP®). It provides a gathering place for Cisco certified professionals to share questions, suggestions,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

https://t.me/learningnets

Your Training Curriculum

This topic presents the training curriculum for this course.



https://t.me/learningnets

https://t.me/learningnets

Module 1

VPN Technologies

Overview

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

Module Objectives

Upon completing this module, you will be able to describe the VPN technologies used in the service provider environment and the MPLS VPN peer-to-peer. You will be able to meet these objectives:

- Explain the concept of VPNs and the VPN terminology
- Explain the MPLS VPN architecture, route information propagation, RDs, RTs, and virtual routing tables

https://t.me/learningnets

https://t.me/learningnets

Lesson 1

Introducing VPNs

Overview

This lesson introduces the main characteristics of VPNs. This lesson explains the concept of VPNs, the count advantages of VPNs, and the VPN terminology that is also used by the Multiprotocol Label Switching (MPLS) VPN architecture. The lesson also explains the differences between the overlay and peer-to-peer VPN models, how they are implemented, and the benefits and drawbacks of each implementation.

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

Objectives

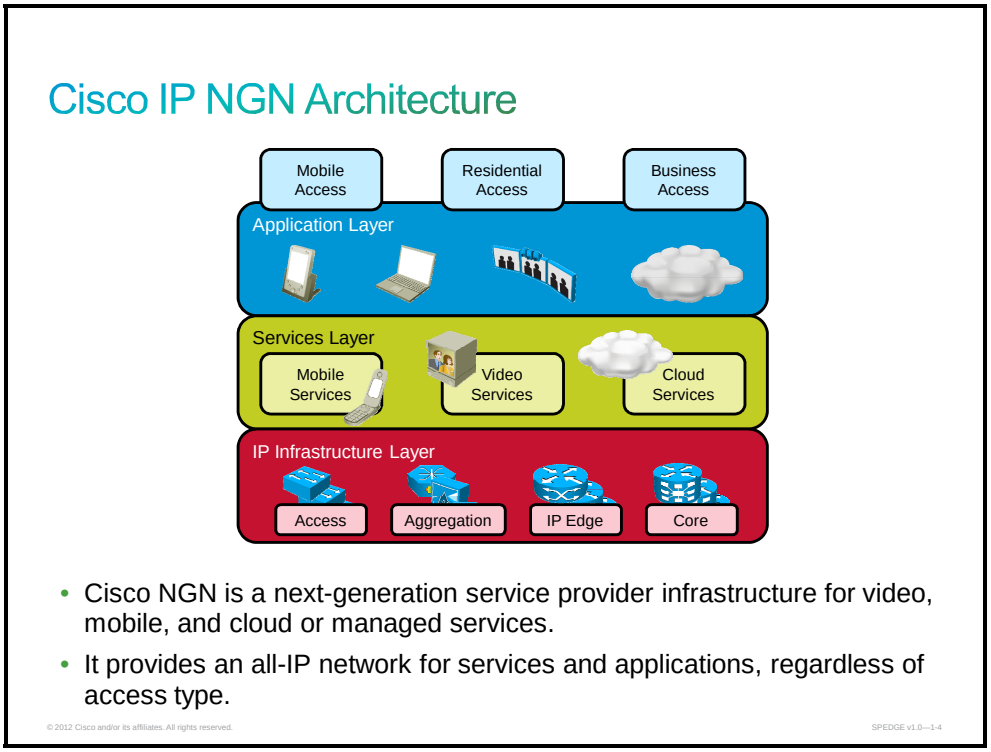
Upon completing this lesson, you will be able to explain the concept of VPNs and the VPN terminology. You will be able to meet these objectives:

- Describe the concept of VPNs and the reasons why VPNs were introduced
- Describe VPN implementation models, and list benefits and drawbacks of VPNs

https://t.me/learningnets

VPN Concept

This topic describes the concept of VPNs and the reasons why VPNs were introduced.



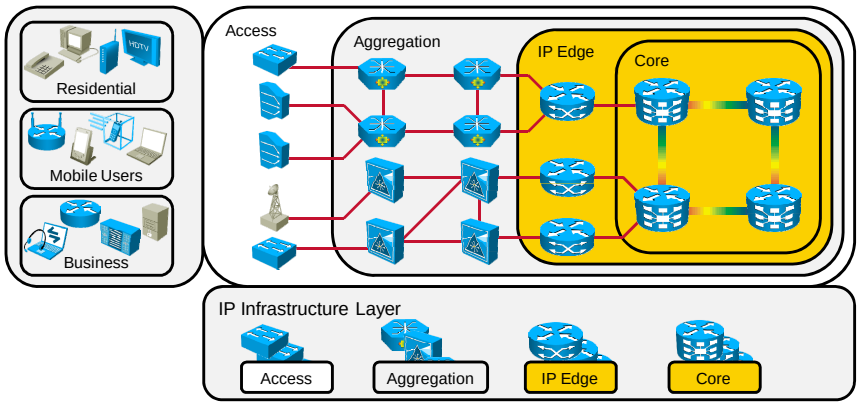
The Cisco IP Next-Generation Network (NGN) architecture enables service providers to start developing fixed and mobile convergence starting with the transport in the access, aggregation, and core networks. The NGN targets service providers with an existing centralized wireline services edge network that are willing to maintain and evolve this network layer as part of their services, network, and organizational evolution.

The NGN architecture provides a flexible, comprehensive, and generic framework that is structured around the most common layers in service provider networks: customer premises, access network, aggregation network, edge network, core network, network management, and network admission layers. The access, aggregation, and core layers are used for transport of mobile, video, and cloud or managed services.

The general idea of the Cisco IP NGN is to provide all-IP transport for all services and applications, regardless of access type. IP infrastructure, service, and application layers are separated in NGNs, thus enabling the addition of new services and applications without any changes in the transport network.

https://t.me/learningnets

Cisco IP NGN Infrastructure Layer



- VPNs relay on the IP edge and core parts of the IP infrastructure layer of the Cisco IP NGN.

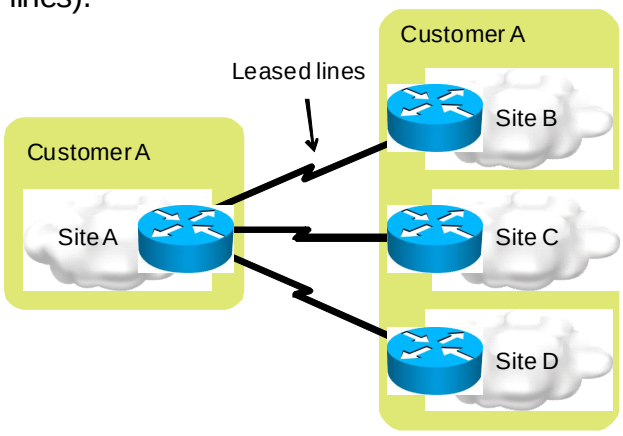
© 2012 Cisco and/or its affiliates. All rights reserved.

VPNs relay on the IP edge and core parts of the IP infrastructure layer of the Cisco IP NGN. Features are primarily configured on the IP edge layer. The IP core layer should be as transparent as possible for scalability reasons. .

https://t.me/learningnets

Traditional Router-Based Networks

Traditional router-based networks connect customer sites through routers connected via dedicated point-to-point links (leased lines).



© 2012 Cisco and/or its affiliates. All rights reserved.

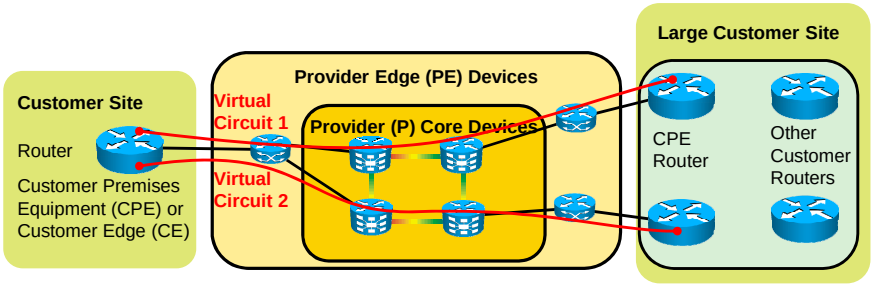
Traditional router-based networks were implemented with dedicated point-to-point links connecting customer sites. The cost of this approach was comparatively high for these reasons:

- The dedicated point-to-point links prevented any form of statistical infrastructure sharing on the service provider side, resulting in high costs for the end user.
- Every link required a dedicated port on a router, resulting in high equipment costs.
- The complexity and cost increased by adding more remote sites.

https://t.me/learningnets

VPN Example

- VPNs replace dedicated point-to-point links with emulated point-to-point links that share common infrastructure.
- Customers use VPNs primarily to reduce their operational costs.



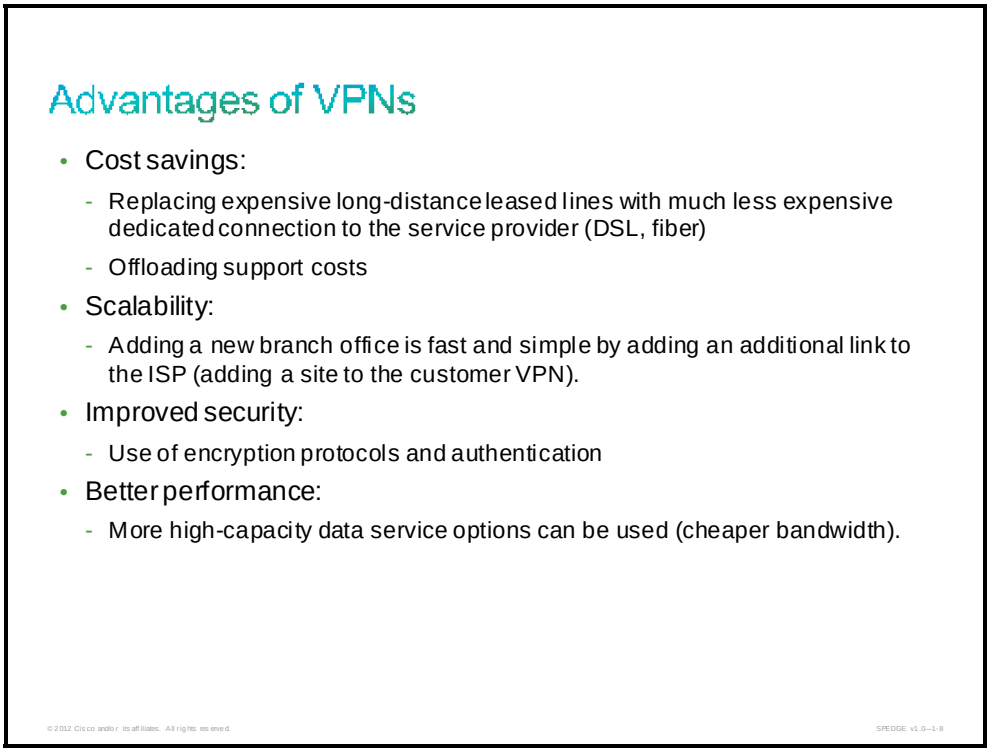
© 2012 Cisco and/or its affiliates. All rights reserved.

VPNs were introduced very early in the history of data communications with technologies such as Frame Relay, which uses virtual circuits (VCs) to establish the end-to-end connection over a shared service provider infrastructure. Although Frame Relay is sometimes considered obsolete, it still shares these basic benefits with modern VPNs:

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

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

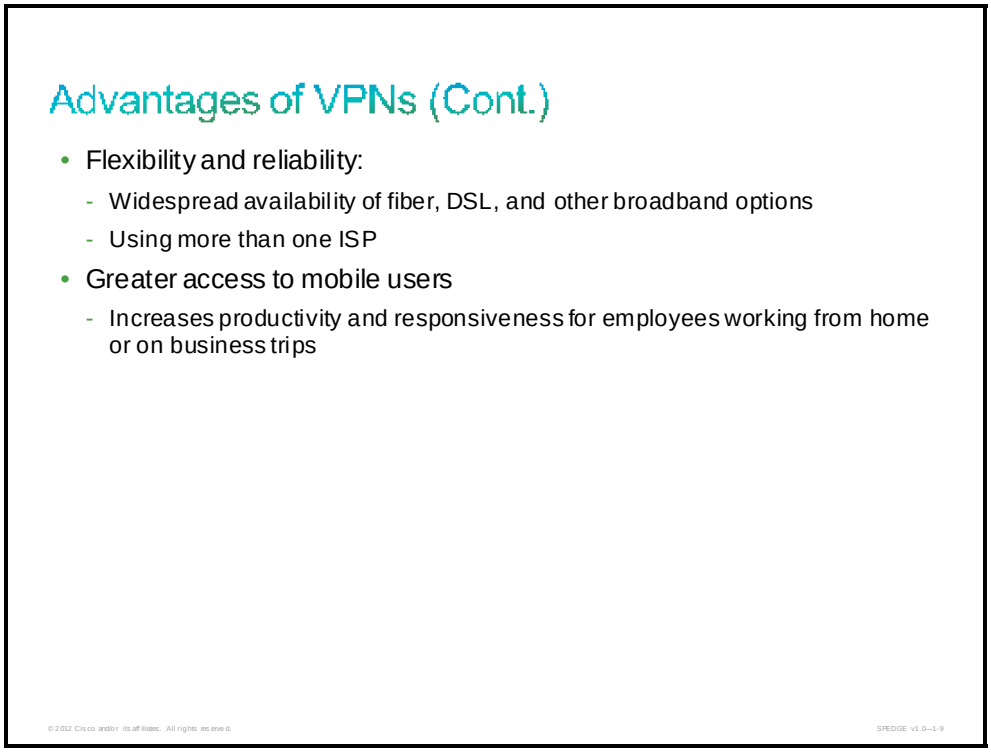
https://t.me/learningnets



VPNs give an organization the advantage of creating secure channels of communication while at the same time reducing costs, improving security, increasing performance, and providing greater access to remote users:

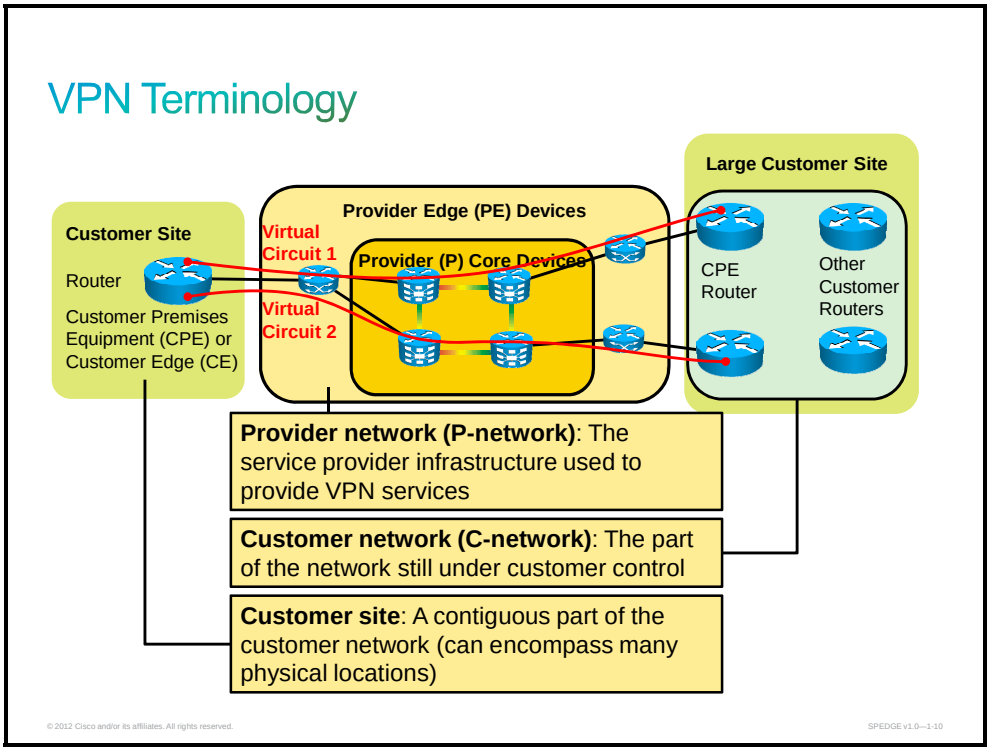
- Cost savings: Dedicated circuits (leased lines) are quite expensive, so replacing leased lines with a much less expensive dedicated connection to the service provider can significantly decrease costs.
- Scalability: A company with two branch offices can deploy just one dedicated line to connect the two locations. If a third branch office needs to come online, two additional lines will be required to directly connect that location to the other two. However, by adding more branch offices to the network, the number of leased lines increases dramatically (four branch offices require six lines for full connectivity; five offices require ten lines, and so on). VPNs avoid this problem by simply adopting one link to ISP per branch office.
- Improved security: The use of encryption protocols and authentication helps secure the data that is traveling over the VPN channel.
- Better performance: VPNs also provide greater bandwidth, thus allowing better performance.

https://t.me/learningnets



- Flexibility and reliability: The widespread availability of fiber, DSL, and other broadband options gives enterprises multiple ways to securely interconnect over a VPN. VPNs can also improve reliability by using data services from several independent ISPs at the same time (having redundant solutions).
- Greater access to mobile users: Many workers can work from home or spend a significant amount of time on business trips. By adopting VPN solutions, they are able to connect from anywhere to company servers to access email and data, thus increasing productivity and responsiveness.

https://t.me/learningnets



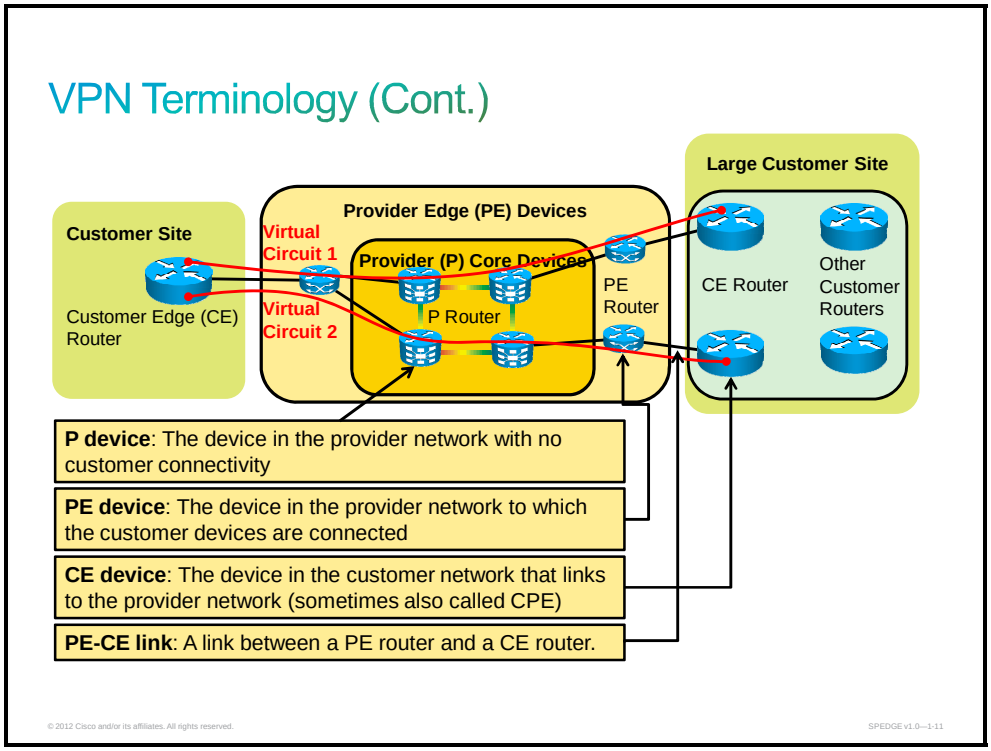
There are many conceptual models and terminologies that describe various VPN technologies and implementations. The terminology is generic enough to cover nearly any VPN technology or implementation and is thus extremely versatile.

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

- **Provider network (P-network):** The common infrastructure that the service provider uses to offer VPN services to customers. Service provider devices to which the customer edge (CE) routers were directly attached were called provider edge (PE) routers. In addition, the service provider network might consist of devices used for forwarding data in the service provider backbone called provider routers (P routers).
- **Customer network (C-network):** The part of the overall customer network that is still exclusively under customer control. It consists of the routers at the various customer sites. The routers connecting the sites of individual customers to the service provider network are called CE routers.
- **Customer sites:** These are contiguous parts of the C-network.

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

https://t.me/learningnets



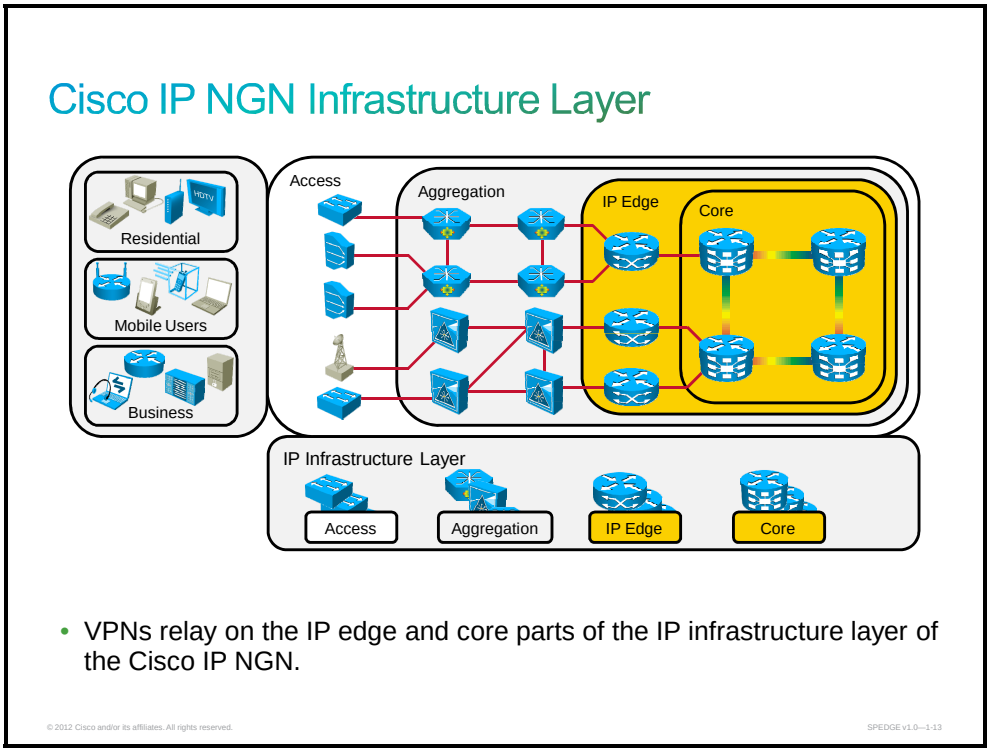
Here is a description of the devices that enable the overall VPN solution. The devices are named based on their position in the network:

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

https://t.me/learningnets

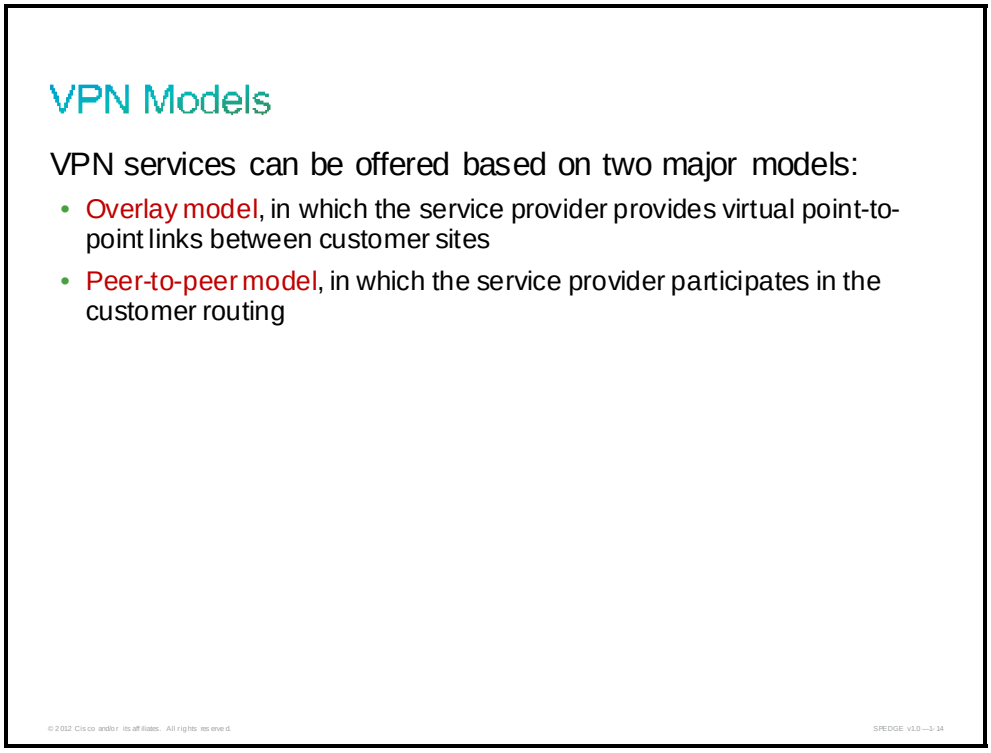
VPN Models

This topic describes VPN implementation models, and lists benefits and drawbacks of VPNs.



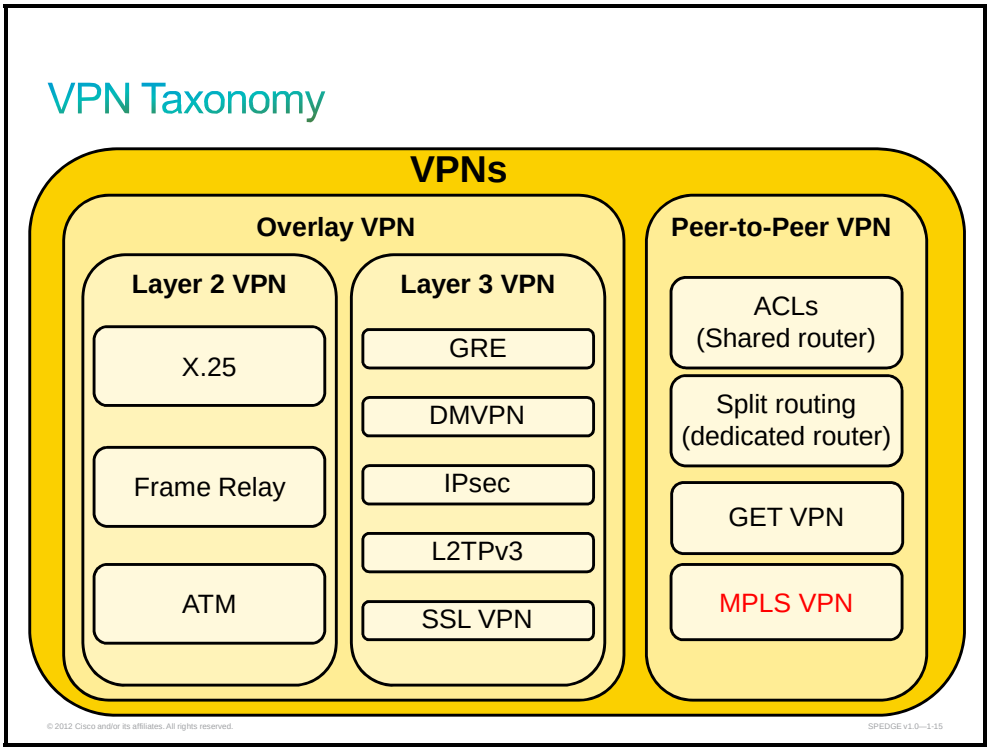
All VPN implementation models relay on the IP edge and core parts of the IP infrastructure layer of the Cisco IP NGN.

https://t.me/learningnets



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

https://t.me/learningnets



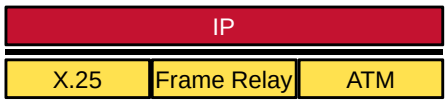
VPNs allow you to use the shared infrastructure of a service provider to implement your private networks. There are basically these two implementation models:

- Overlay VPNs
  - Layer 2, including technologies such as X.25, Frame Relay, and ATM
  - Layer 3, including Generic Routing Encapsulation (GRE), Dynamic Multipoint VPN (DMVPN), IPsec, SSL VPN, and Layer 2 Tunneling Protocol (L2TP)
- Peer-to-peer VPNs, implemented with routers and respective filters, separate routers per customer via GET VPN or with MPLS VPN technology, which is covered in greater detail in later lessons.

https://t.me/learningnets

Overlay Layer 2 VPNs

- Layer 2 VPN
  - The service provider establishes Layer 2 VCs between customer sites.
  - The customer is responsible for all higher layers.



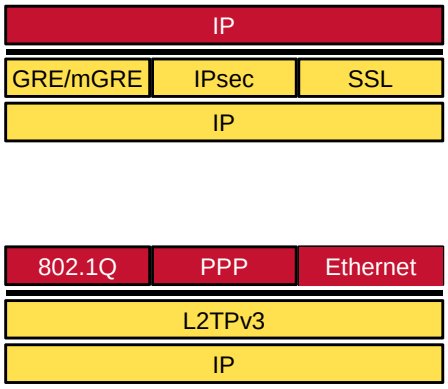
© 2012 Cisco and/or its affiliates. All rights reserved.

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

https://t.me/learningnets

Overlay Layer 3 VPN

- The service provider infrastructure appears as point-to-point links to the customer.
- The service provider does not see customer routes and is responsible only for providing the point-to-point transport of customer data.
- Layer 3 VPN – IP tunneling
  - Routing protocols run directly between customer routers.
  - GRE is simple (and quicker).
  - IPsec provides authentication and security.
- Layer 2 VPN – Layer 2 forwarding
  - Transparent tunneling of Layer 2 over IP



© 2012 Cisco and/or its affiliates. All rights reserved.

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

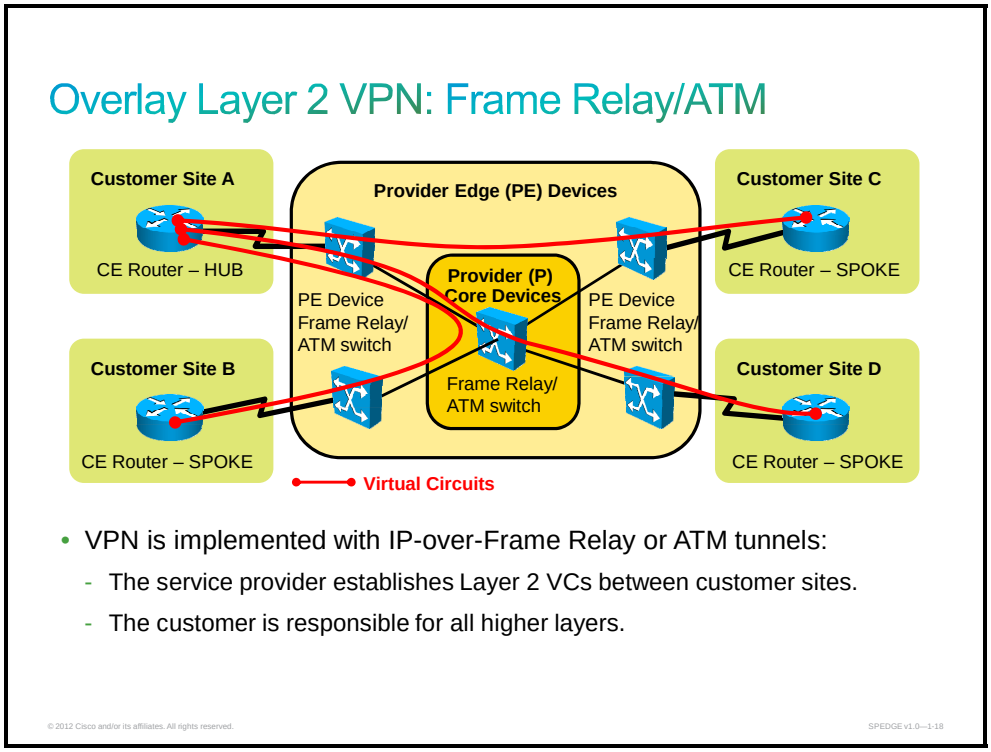
Whatever the business reasons behind it, Layer 3 VPN implementations over the IP backbone always involve tunneling—encapsulation of protocol units at a certain layer of the Open Systems Interconnection (OSI) reference model into protocol units at the same or a higher layer of the OSI model.

Two well-known tunneling technologies are IP Security (IPsec) and GRE. GRE is fast and simple to implement and supports multiple routed protocols, but it provides no security and is thus unsuitable for deployment over the Internet. An alternative tunneling technology is IPsec, which provides network layer authentication and optional encryption to make data transfer over the Internet secure. IPsec supports only the IP routed protocol. SSL is the latest method to make authentication and encryption of data transferred over the Internet secure. It is a remote access solution that replaces IPsec clients and is firewall-friendly (uses SSL as the transport).

Layer 2 Tunnel Protocol Version 3 (L2TPv3) is capable of tunneling any Layer 2 payload over L2TP.

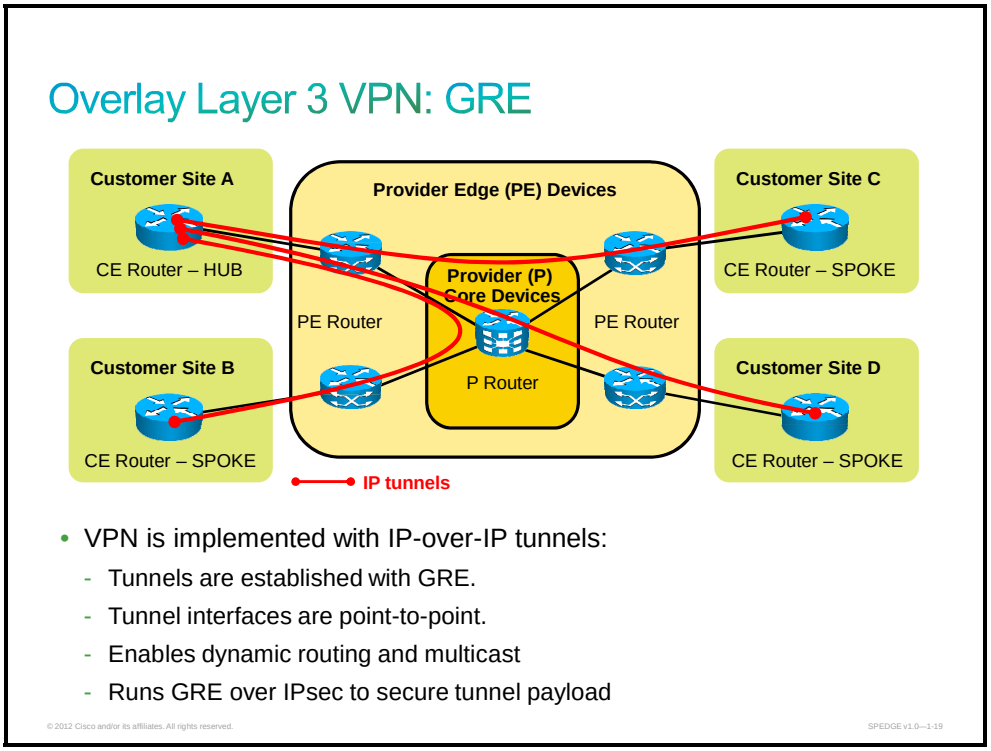
https://t.me/learningnets

The figure shows a typical Layer 2 overlay VPN implemented by a Frame Relay network.



The customer needs to connect three sites to Site A (central site, hub site) and orders connectivity between Site A (hub) and Site B (spoke), between Site A and Site C (spoke), and between Site A and Site D (spoke). The service provider implements this request by providing three permanent virtual circuits (PVCs) across the Frame Relay network, thus enabling Layer 2 connectivity between hub and spoke sites. Note that spoke-to-spoke traffic has to go through the hub site.

https://t.me/learningnets

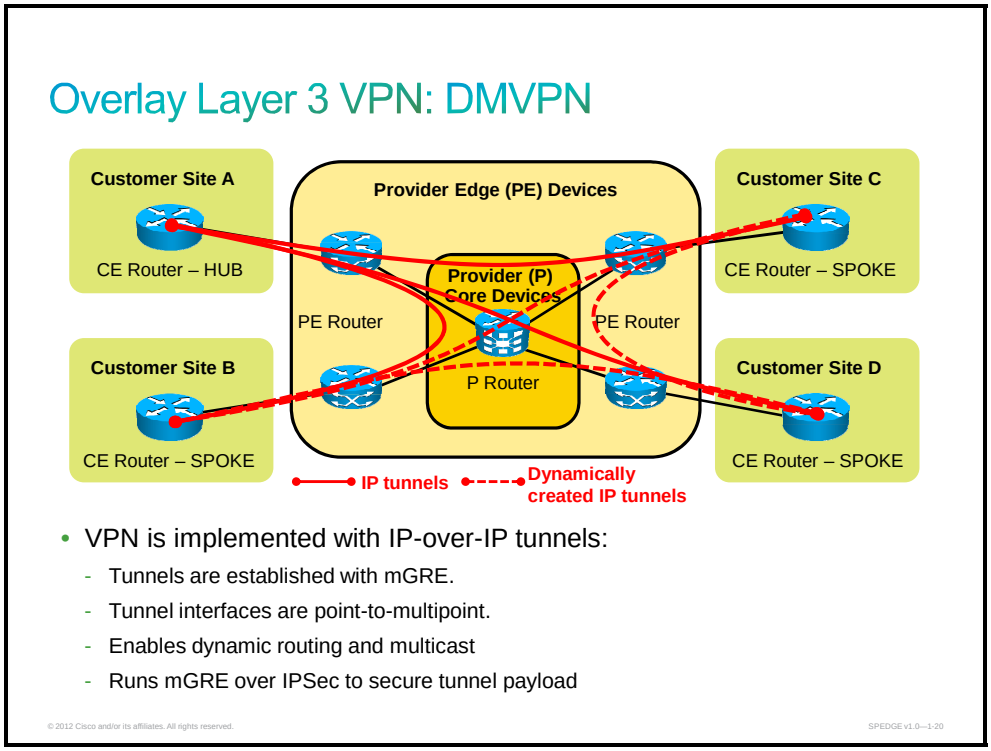


The figure presents the same scenario as the previous figure (implemented by a Frame Relay network). The difference is that, in this case, Layer 3 connectivity is provided between hub and spoke sites by using GRE point-to-point tunnels. The customer needs to connect three sites to Site A (central site, hub site) and orders connectivity between Site A (hub) and Site B (spoke), between Site A and Site C (spoke), and between Site A and Site D (spoke). The service provider implements IP connectivity over its network. Note that spoke-to-spoke traffic has to go through the hub site.

The GRE is a multiprotocol-capable transport protocol (IPv4, IPv6, MPLS, and so on) and enables dynamic routing and multicast over the tunnels.

To secure the tunnel payload, you have to run GRE over IPsec.

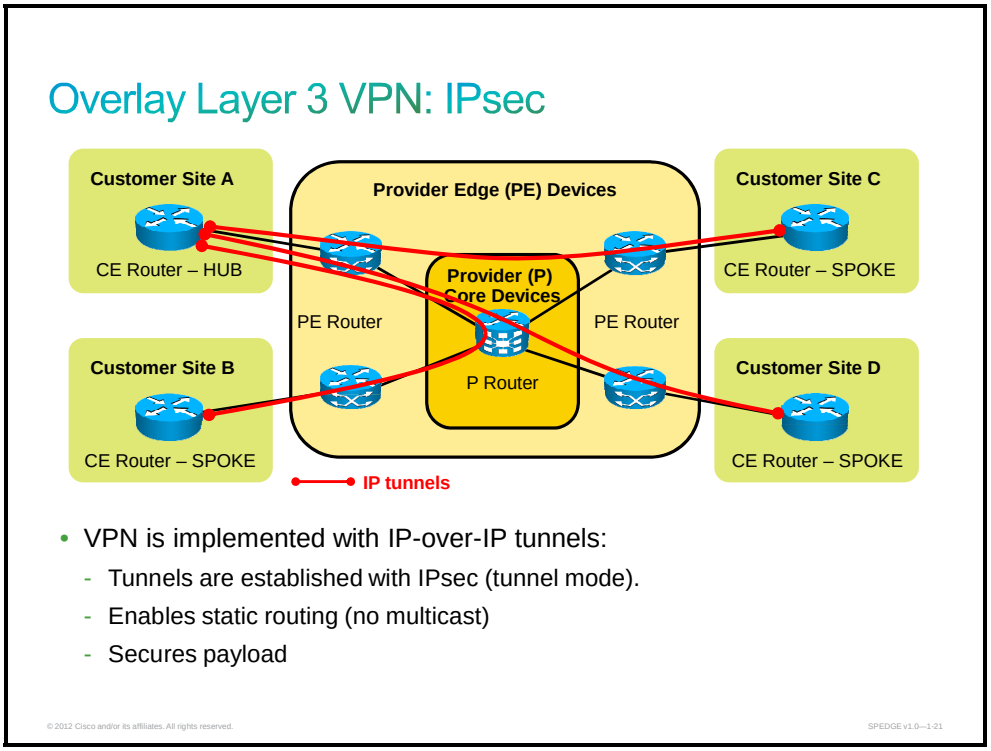
https://t.me/learningnets



In this DMVPN scenario, point-to-multipoint GRE (mGRE) tunnels are used. The customer needs to connect three sites to Site A (central site, hub site) and orders connectivity between Site A (hub) and Site B (spoke), between Site A and Site C (spoke), and between Site A and Site D (spoke). The service provider implements IP connectivity over its network. Note that in this DMVPN scenario, spoke-to-spoke traffic can flow directly by dynamically establishing GRE tunnels between spokes. To secure the tunnel payload, you have to run mGRE over IPsec.

The GRE is a multiprotocol-capable transport protocol (IPv4, IPv6, MPLS, and so on) and enables dynamic routing and multicast over the tunnels.

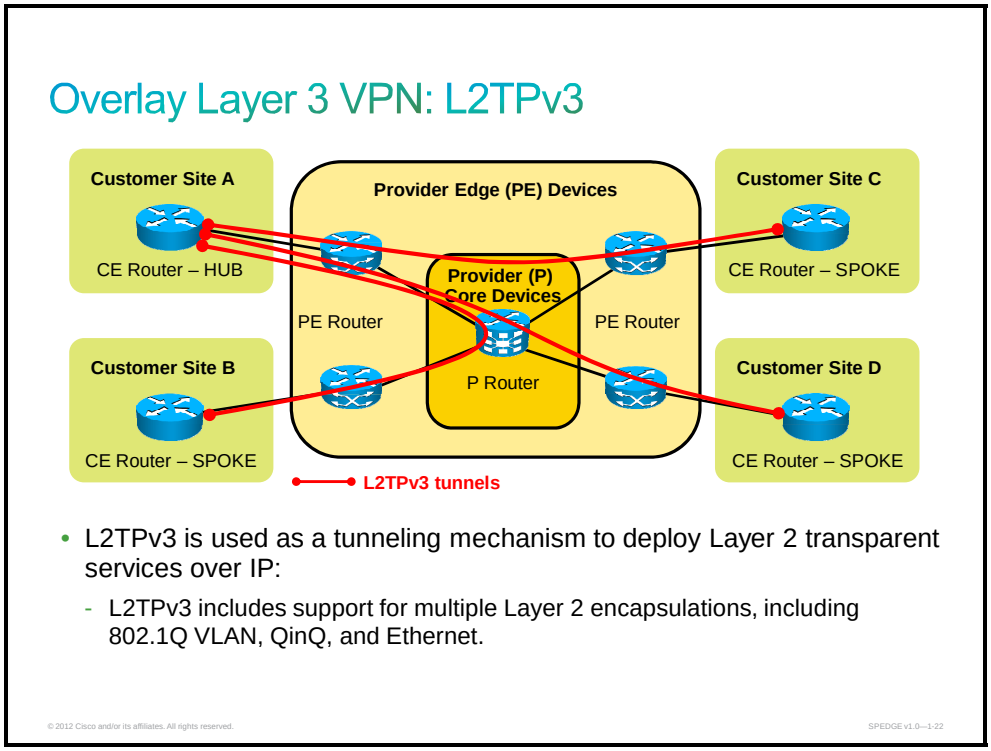
https://t.me/learningnets



IPsec provides network layer authentication and optional encryption to make data transfer over the Internet secure. This is achieved by creating IP-over-IP tunnels and securing the payload. The limitation of the IPsec tunnels is that they do not offer multicast functionality, instead providing static routing only.

The usage of IPsec, that is, securing the payload, is usually used in securing GRE and mGRE tunnels.

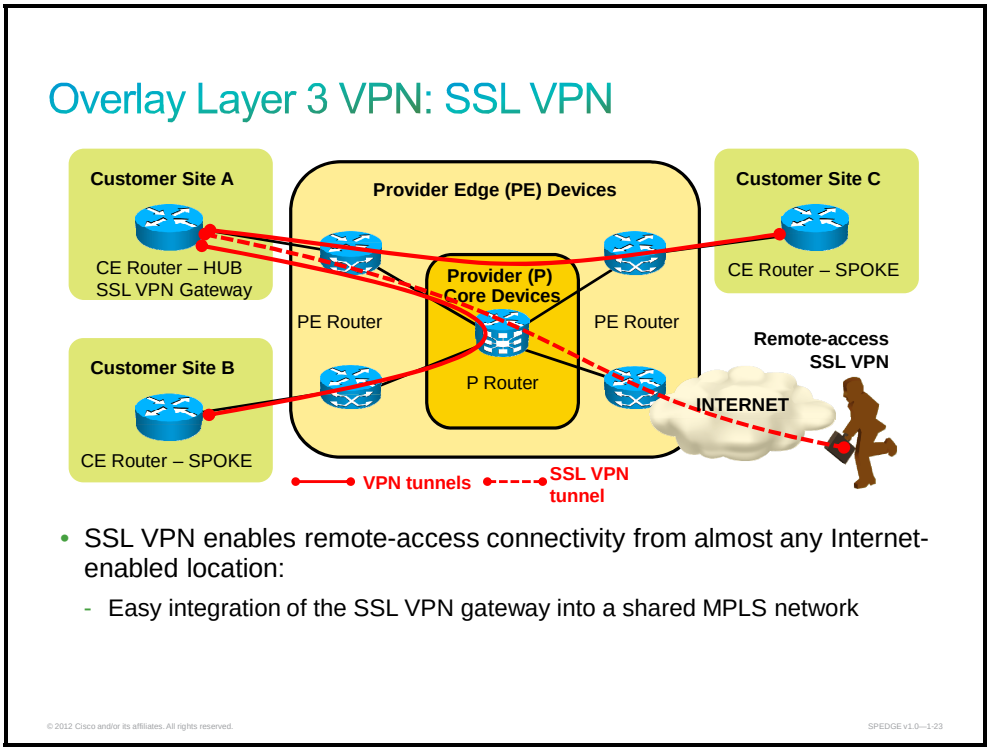
https://t.me/learningnets



Layer 2 Tunnel Protocol version 3 (L2TPv3) is capable of tunneling any Layer 2 payload over L2TP. Specifically, L2TPv3 defines the L2TP protocol for tunneling Layer 2 payloads over an IP core network using Layer 2 VPNs. The benefits of this feature include the following:

- L2TPv3 simplifies deployment of VPNs
- L2TPv3 does not require MPLS
- L2TPv3 supports Layer 2 tunneling over IP for any payload

https://t.me/learningnets

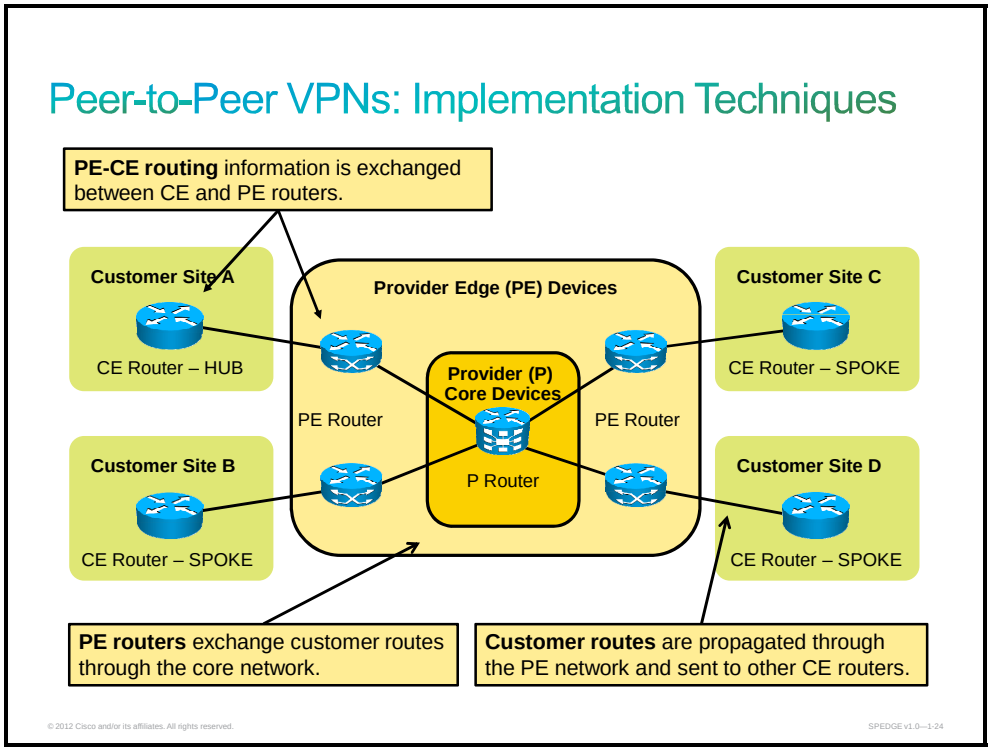


Secure Sockets Layer (SSL) is the method to achieve secure authentication and encryption of the data transfer over the Internet. It is a remote access solution that replaces IPsec clients and is firewall-friendly (uses SSL as the transport). It runs in three operational models:

- Clientless, providing access to web servers behind the firewall
- Thin client, providing port forwarding via a Java applet
- Full tunnel with SSL VPN client

It is possible to integrate an SSL VPN gateway into an MPLS VPN network.

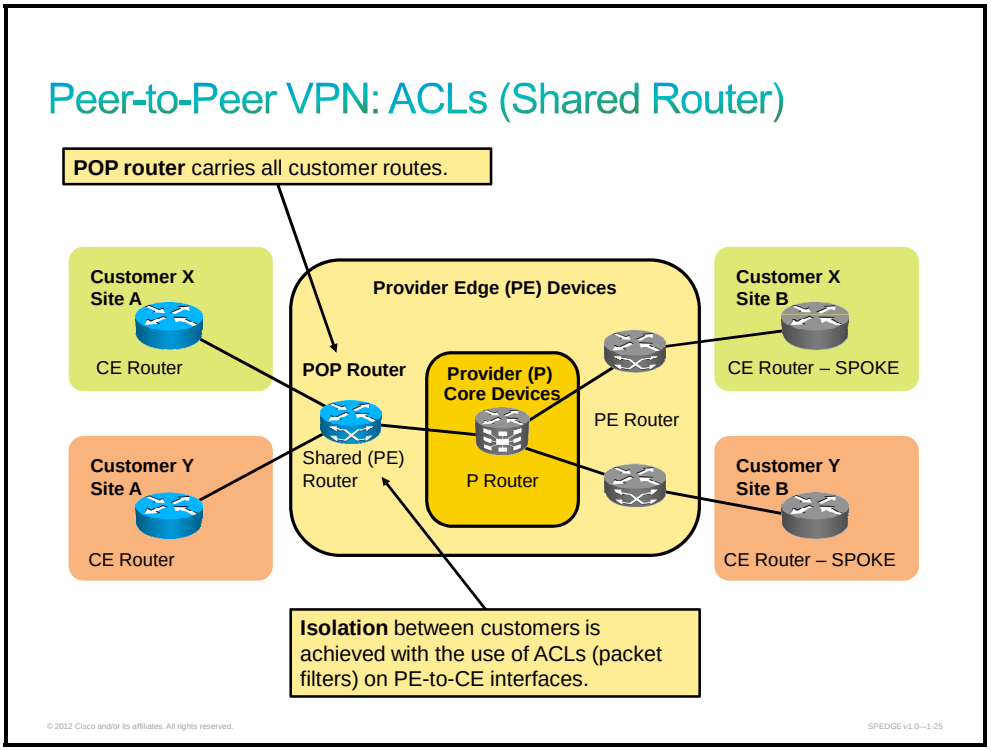
https://t.me/learningnets



The point-to-point overlay VPN model has a number of drawbacks, most significantly the need for customers to establish point-to-point links or virtual circuits (VCs) between sites. The formula to calculate how many point-to-point links or VCs are needed is ([n]*[n-1])/2, where n is the number of sites to be connected. For example, if a customer wants to have a full mesh between 10 sites, it would need 10*9/2=45 point-to-point links. This would certainly be a scalability issue.

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

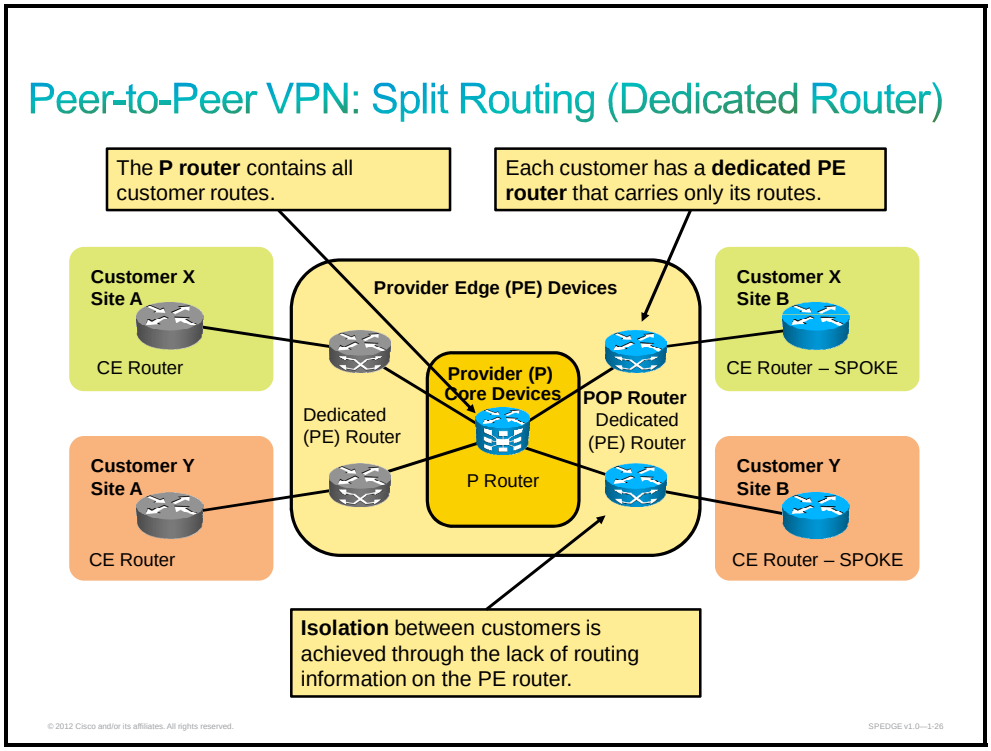
https://t.me/learningnets



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

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

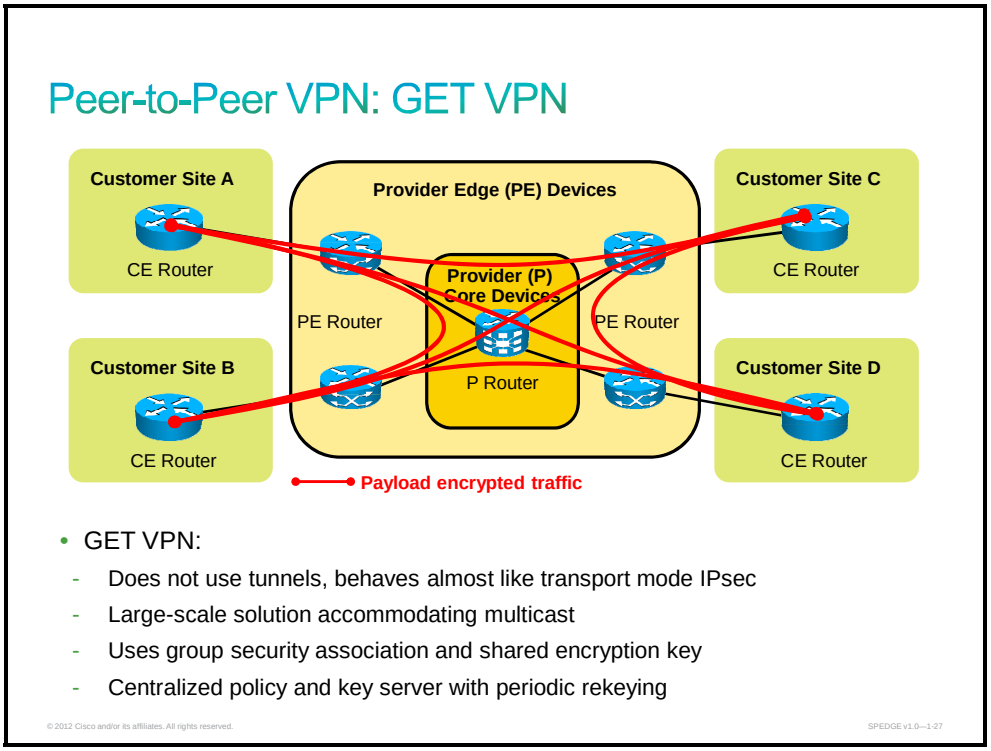
https://t.me/learningnets



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

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

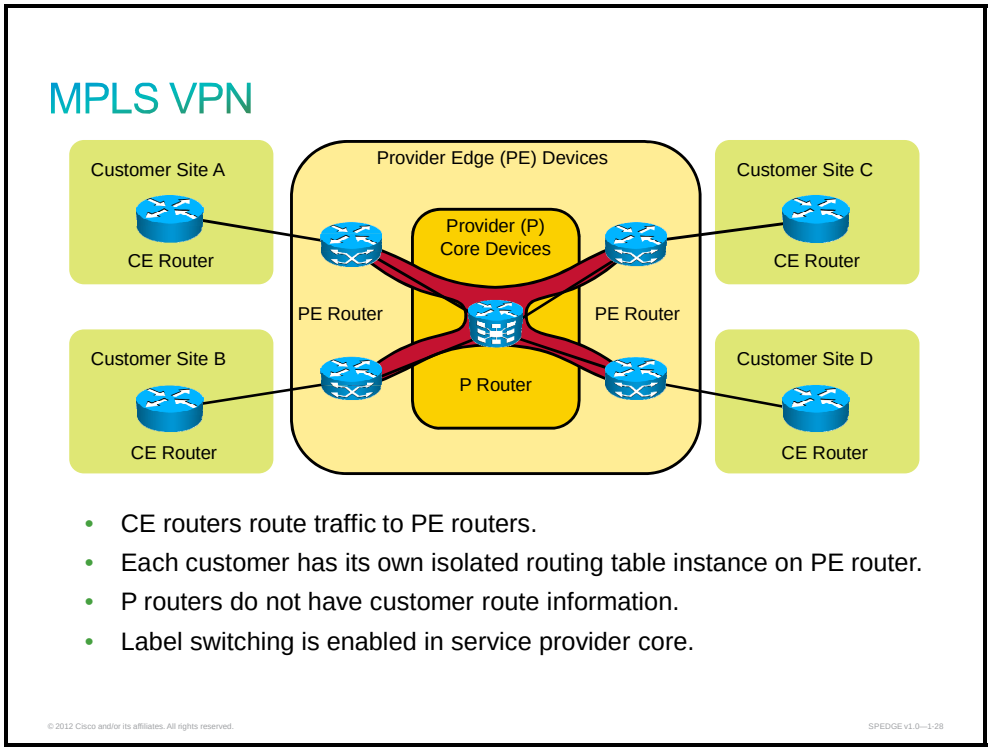
https://t.me/learningnets



GET VPN is a tunnel-less VPN technology that provides end-to-end security for network traffic in a native mode and maintains the fully meshed topology. GET VPN preserves the original source and destination IP addresses information in the header of the encrypted packet for optimal routing (like transport mode IPsec). Hence, it is largely suited for an enterprise running over a private MPLS and IP-based core network. It is also better suited to encrypt multicast traffic. GET VPN uses Group Domain of Interpretation (GDOI) as the keying protocol and IPsec for encryption. Some of the advantages of GET VPN are as follows:

- Provides high scalability to any meshed topology and eliminates the need for complex peer-to-peer security associations.
- For MPLS networks, maintains network intelligence (such as full-mesh connectivity, natural routing path, and quality of service [QoS]). Grants easy membership control with centralized key servers.
- Helps ensure low latency and jitter by enabling full-time, direct communications between sites without requiring transport through a central hub.
- Allows replication of the packets after encryption. This allows the multicast traffic to be replicated at the core, thereby reducing the load and bandwidth requirement on the customer premises equipment (CPE).
- IP address preservation enables encrypted packets to carry the original source and destination IP addresses in the outer IP header rather than replacing them with tunnel endpoint addresses. This technique is known as IPsec tunnel mode with address preservation. Some of the IP header parameters are also preserved. Many network features like routing, basic firewall, QoS, and traffic management work based on the information contained in the IP header. Since the IP header is persevered, all the network features will work as before. This eliminates many issues associated with deploying point-to-point encryption in a core network.

https://t.me/learningnets



In the MPLS VPN model, the best features of the overlay and point-to-point models are implemented.

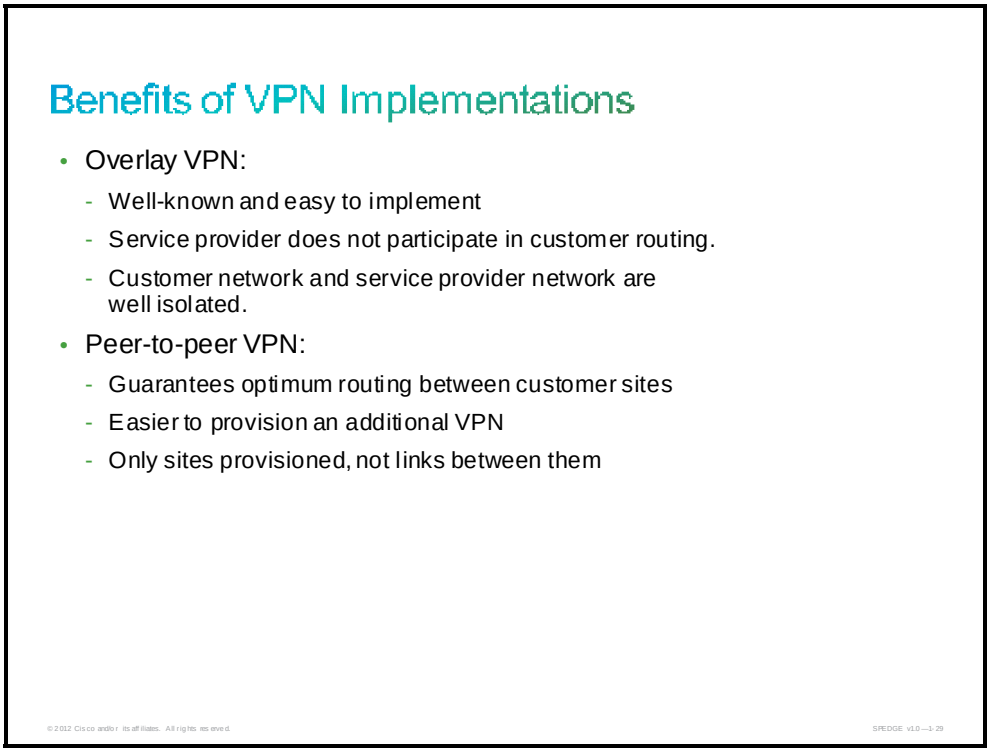
An MPLS-enabled core and edge network provides a very fast and efficient data switching environment based on MPLS labels.

PE routers exchange routing information with customer CE routers and use separate isolated routing tables for each customer. Special routing protocol contexts are used for route exchange between PE and CE routers.

Routes are then exchanged between PE devices using the Multiprotocol BGP (MP-BGP) routing algorithm.

For scalability reasons, service provider core routers do not have any customer routing information. PE routers label packets with MPLS labels and P routers use these labels for fast label-switching packets.

https://t.me/learningnets



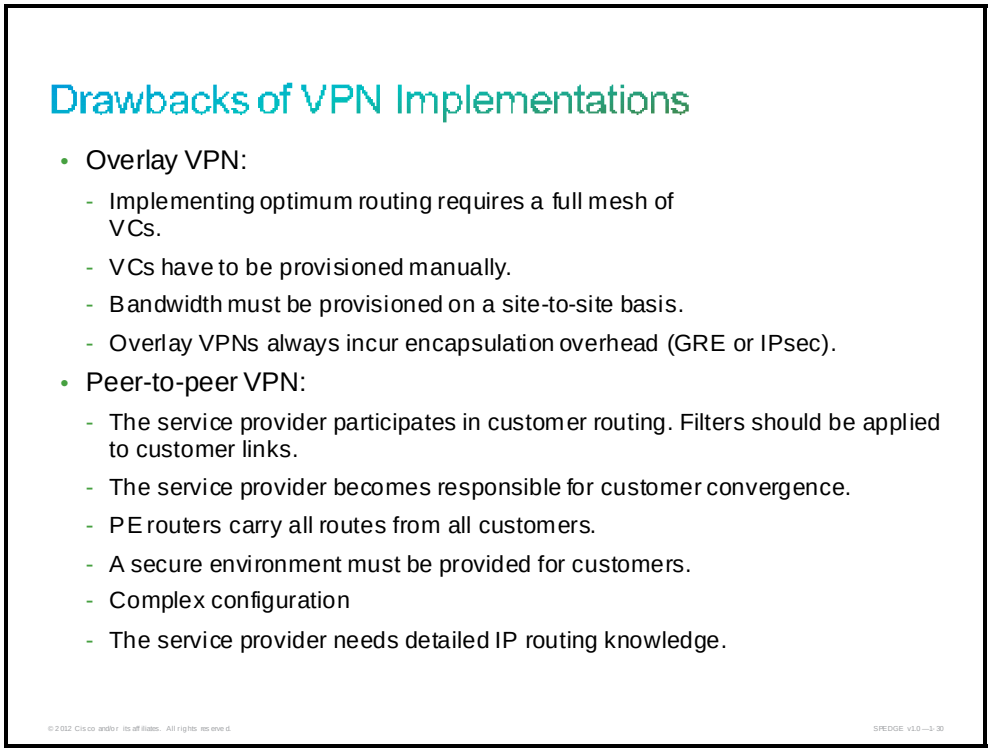
Each VPN model has a number of benefits. For example, overlay VPNs have these advantages:

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

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

- They provide optimum routing between customer sites without any special design or configuration effort.
- They offer easy provisioning of additional VPNs or customer sites, because the service provider provisions only individual sites, not the links between individual customer sites

https://t.me/learningnets



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

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

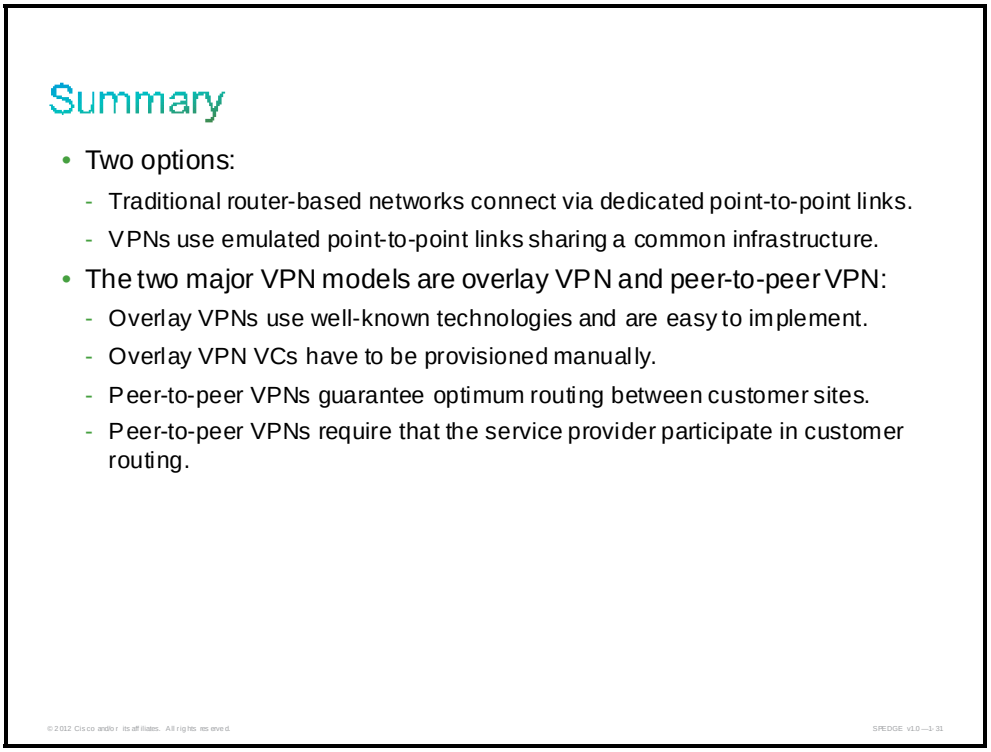
The major drawbacks of peer-to-peer VPNs arise from service provider involvement in customer routing, such as these disadvantages:

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

https://t.me/learningnets

Summary

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



https://t.me/learningnets

Lesson 2

Introducing MPLS VPNs

Overview

This lesson explains the Multiprotocol Label Switching (MPLS) VPN architecture. It is important to understand how the MPLS VPN architecture is structured, what the components of that architecture are, and how the components are used. This knowledge will help later, when you begin to look at design issues and configuration parameters. The lesson offers address and routing perspectives from the customer and service provider side, and it discusses how routing tables appear on provider edge (PE) routers. This lesson also explains how forwarding across an MPLS VPN backbone occurs, identifies how labels get propagated, and explains the effects of summarization in the core.

Objectives

Upon completion of this lesson, you will be able to understand MPLS VPNs. You will be able to meet these objectives:

- Explain the MPLS VPN architecture, RDs, RTs, and virtual routing tables
- Describe end-to-end routing update flow
- Describe VPN label propagation between PE routers and the MPLS VPN end-to-end forwarding mechanism

https://t.me/learningnets

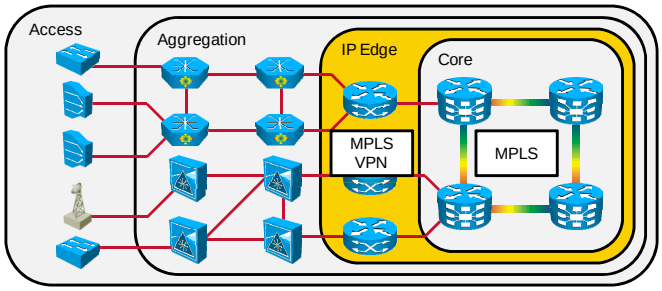
MPLS VPN Architecture

This topic explains the MPLS VPN architecture, route distinguishers (RDs), route targets (RTs), and virtual routing tables.

MPLS VPN Architecture

An MPLS VPN combines the best features of an overlay VPN and a peer-to-peer VPN:

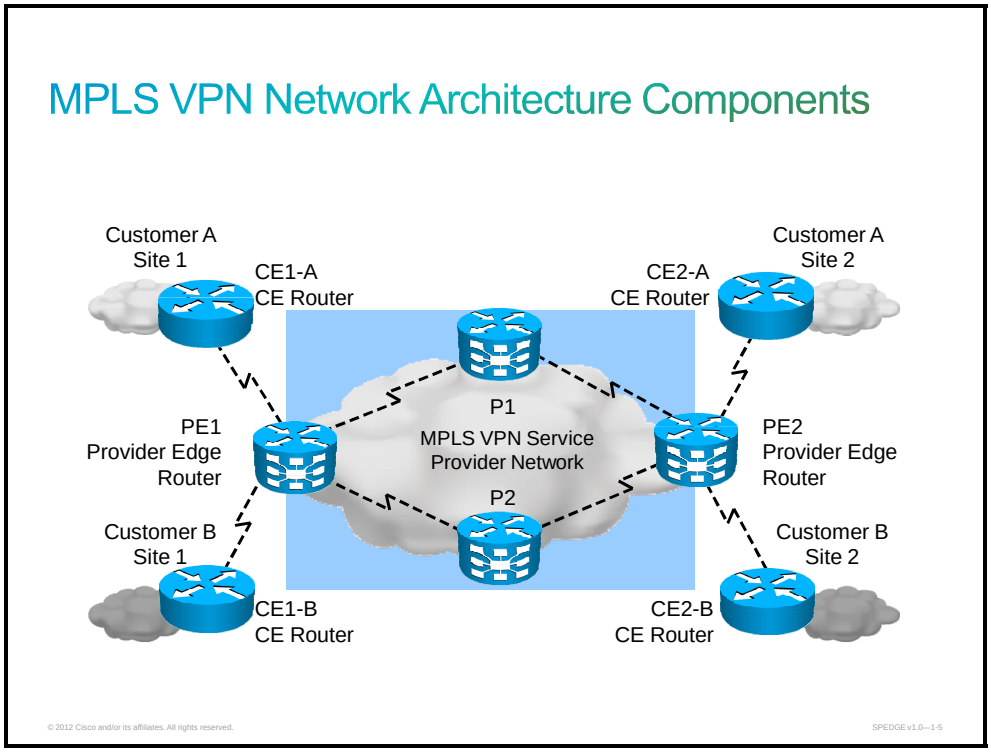
- PE routers participate in customer routing, guaranteeing optimum routing between sites and easy provisioning.
- PE routers carry a separate set of routes for each customer (similar to the dedicated PE router approach).
- Customers can use overlapping addresses.



© 2012 Cisco and/or its affiliates. All rights reserved.

In the MPLS VPN architecture, the edge routers carry customer routing information, providing optimal routing for traffic that belongs to the customer for intersite traffic. The MPLS-based VPN model also accommodates customers who use overlapping address spaces, unlike the traditional peer-to-peer model, in which optimal routing of customer traffic required the provider to assign IP addresses to each of its customers (or the customer to implement Network Address Translation [NAT]) to avoid overlapping address spaces. MPLS VPN is an implementation of the peer-to-peer model; the MPLS VPN backbone and customer sites exchange Layer 3 customer routing information, and data is forwarded between customer sites using the MPLS-enabled service provider IP backbone.

https://t.me/learningnets



The MPLS VPN domain, like the traditional VPN, consists of the customer network and the provider network. The MPLS VPN model is very similar to the dedicated provider edge (PE) router model in a peer-to-peer VPN implementation. However, instead of deploying a dedicated PE router per customer, customer traffic is isolated on the same PE router that provides connectivity into the service provider network for multiple customers.

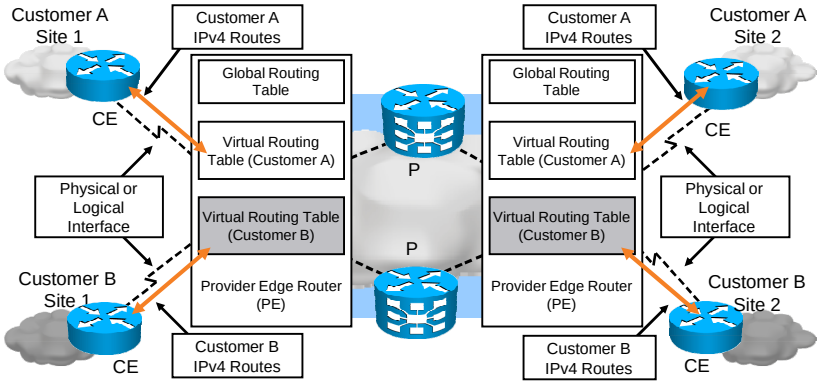
The main components of MPLS VPN architecture are as follows:

- **Customer network**: Usually a customer-controlled domain consisting of devices or routers spanning multiple sites that belong to the customer.
- **Customer edge (CE) routers**: Routers in the customer network that interface with the service provider network.
- **Provider network**: The provider-controlled domain consisting of PE and provider core routers that connect sites belonging to the customer on a shared infrastructure. The provider network controls the traffic routing between sites belonging to a customer, along with customer traffic isolation.
- **PE routers**: Routers in the provider network that interface or connect to the customer edge routers in the customer network.
- **P routers**: Routers in the core of the provider network that interface with either other provider core routers or provider edge routers.

Note	In an MPLS VPN implementation, the PE router is the edge label switch router (edge LSR).

https://t.me/learningnets

PE Router Architecture



A PE router in an MPLS VPN uses virtual routing tables to implement the functionality of customer-dedicated PE routers.

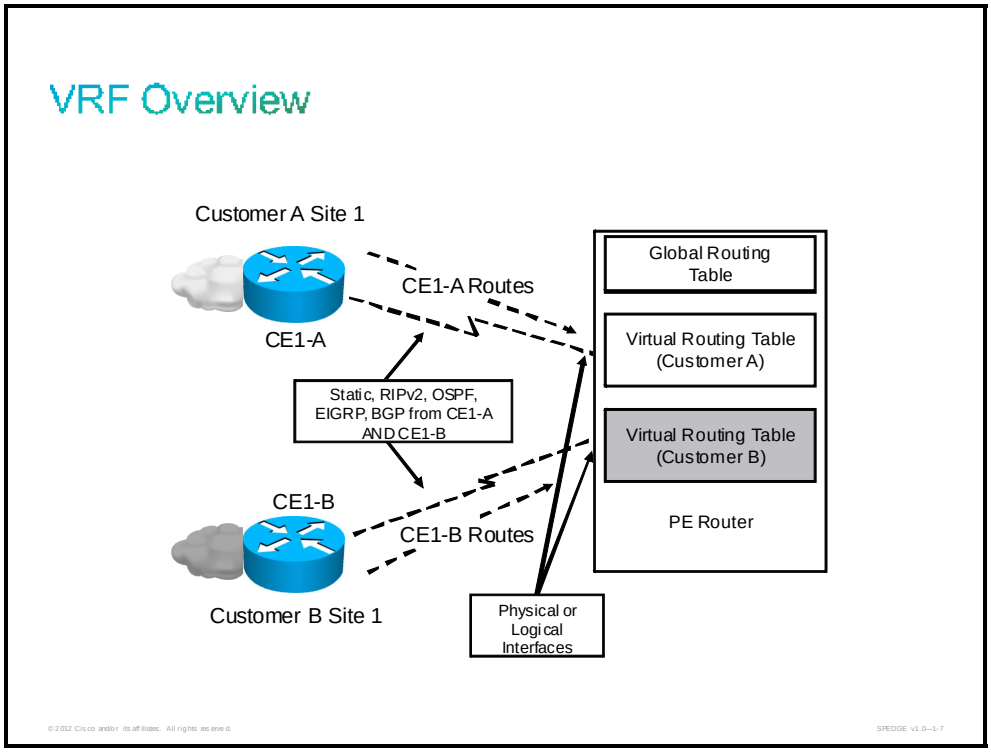
© 2012 Cisco and/or its affiliates. All rights reserved.

An MPLS VPN implementation is very similar to a dedicated router peer-to-peer model implementation. From the perspective of a CE router, only IPv4 updates, as well as data, are forwarded to the PE router. The CE router does not need any specific configuration to enable it to be a part of an MPLS VPN domain. The only requirement on the CE router is a routing protocol (or a static default route) that enables the router to exchange IPv4 routing information with the connected PE router.

In the MPLS VPN implementation, the PE router performs multiple functions. The PE router must first be capable of isolating customer traffic if more than one customer is connected to the PE router. Each customer, therefore, is assigned an independent routing table (virtual routing table or virtual routing and forwarding [VRF] table) similar to a dedicated PE router in the initial peer-to-peer discussion. Routing across the service provider backbone is performed using a routing process in the global routing table. P routers provide label switching between PE routers and are unaware of VPN routes. CE routers in the customer network are not aware of the P routers, and thus the internal topology of the service provider network is transparent to the customer. The figure shows the functionality of the PE router.

The P routers are responsible only for label switching of packets. They do not carry VPN routes and do not participate in MPLS VPN routing. The PE routers exchange IPv4 routes with connected CE routers using individual routing protocol contexts. To enable scaling the network to a large number of customer VPNs, Multiprotocol Border Gateway Protocol (MP-BGP) is configured between PE routers to carry customer routes.

https://t.me/learningnets

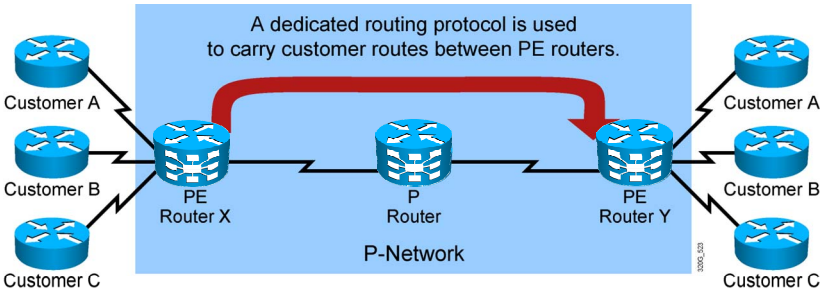


The VRF contains an IP routing table that is analogous to the global IP routing table, a Cisco Express Forwarding table, a list of interfaces that are part of the VRF, and a set of rules defining routing protocol exchange with attached CE routers (routing protocol contexts). In addition, the VRF also contains VPN identifiers and VPN membership information (route distinguisher [RD] and route target [RT] are covered in the next section). The interface that is part of the VRF must support Cisco Express Forwarding switching. The number of interfaces that can be bound to a VRF is limited only by the number of interfaces on the router, and a single interface (logical or physical) can be associated with only one VRF.

The figure shows the function of a VRF on a PE router to implement customer routing isolation. Cisco IOS Software supports various routing protocols and individual routing processes (Open Shortest Path First [OSPF], Enhanced Interior Gateway Routing Protocol [EIGRP], and so on) per router. However, for some routing protocols, such as Routing Information Protocol (RIP) and Border Gateway Protocol (BGP), Cisco IOS Software supports only a single instance of the routing protocol. Therefore, to implement per-VRF routing using protocols that are completely isolated from other VRFs, which might use the same provider edge-customer edge (PE-CE) routing protocols, the concept of routing context was developed.

https://t.me/learningnets

Propagation of Routing Information Across the P-Network



- Run a single routing protocol that will carry all customer routes between PE routers. Use MPLS labels to exchange packets between PE routers.
- P routers do not carry customer routes; the solution is scalable.
- The number of customer routes can be very large. BGP4 is the only routing protocol that can scale to a very large number of routes.
- BGP is used to exchange customer routes directly between PE routers.
- Extend the customer addresses to make them unique.

© 2012 Cisco and/or its affiliates. All rights reserved.

Although virtual routing tables provide isolation between customers, the data from these routing tables still needs to be exchanged between PE routers to enable data transfer between sites attached to different PE routers. Therefore, a routing protocol is needed that will transport all customer routes across the P-network while maintaining the independence of individual customer address spaces.

The best solution to the customer route propagation issue is to run a single routing protocol between PE routers that will exchange all customer routes without the involvement of the P routers. This solution is scalable. Some of the benefits of this approach are as follows:

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

The next design decision to be made is the choice of the routing protocol running between PE routers. Given that the total number of customer routes is expected to be very large, the only well-known protocol with the required scalability is Border Gateway Protocol version 4 (BGP4). In fact, BGP4 is used in the MPLS VPN architecture to transport customer routes directly between PE routers.

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

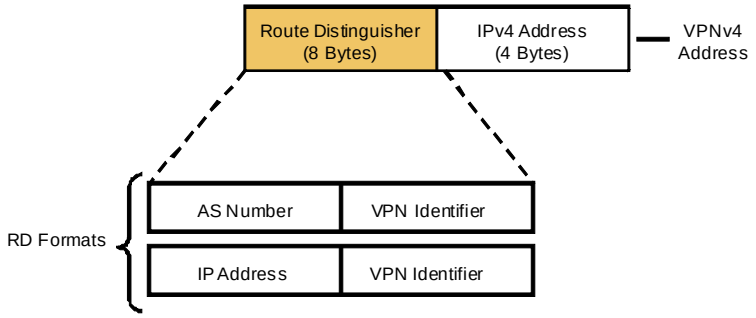
With the deployment of a single routing protocol (that is, BGP4) exchanging all customer routes between PE routers, an important issue arises: how can BGP4 propagate several identical prefixes, belonging to different customers, between PE routers?

The only solution to this dilemma is the expansion of customer IP prefixes with a unique prefix that makes them unique even if they had previously overlapped. A 64-bit prefix called the route distinguisher (RD) is used in MPLS VPNs to convert non-unique 32-bit customer addresses into 96-bit unique addresses that can be transported between PE routers.

https://t.me/learningnets

Route Distinguishers

- In the MPLS VPN backbone, the PE router needs to implement processes that enable overlapping address spaces in connected customer networks.
- The 64-bit route distinguisher is prepended to an IPv4 address to make it globally unique.
- The resulting address is a VPNv4 address.
- VPNv4 addresses are exchanged between PE routers via BGP4.



© 2012 Cisco and/or its affiliates. All rights reserved.

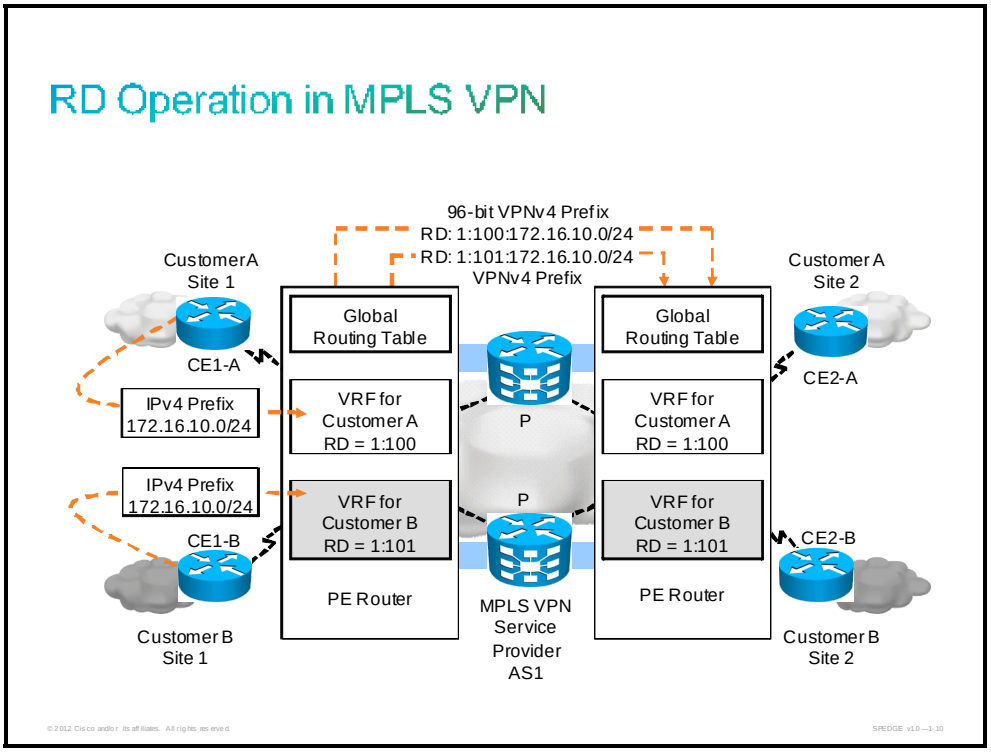
In the MPLS VPN routing model, the PE router provides isolation between customers using VRFs. However, this information must be carried between PE routers to enable data transfer between customer sites via the MPLS VPN backbone. The PE router must be capable of implementing processes that enable overlapping address spaces in connected customer networks. The PE router must also learn these routes from attached customer networks and propagate this information using the shared provider backbone. This is done by the association of an RD per virtual routing table on a PE router.

An RD is a 64-bit unique identifier that is prepended to the 32-bit customer prefix or route learned from a CE router, which makes it a unique 96-bit address that can be transported between the PE routers in the MPLS domain. Thus, a unique RD is configured per VRF on the PE router. The resulting address, which is 96 bits total (32-bit customer prefix plus 64-bit unique identifier or RD), is called a VPNv4 address.

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

The format of an RD is shown in the figure. An RD can be of two formats. If the provider does not have a BGP autonomous system (AS) number, the IP address format can be used, and if the provider does have an AS number, the AS number format can be used.

https://t.me/learningnets



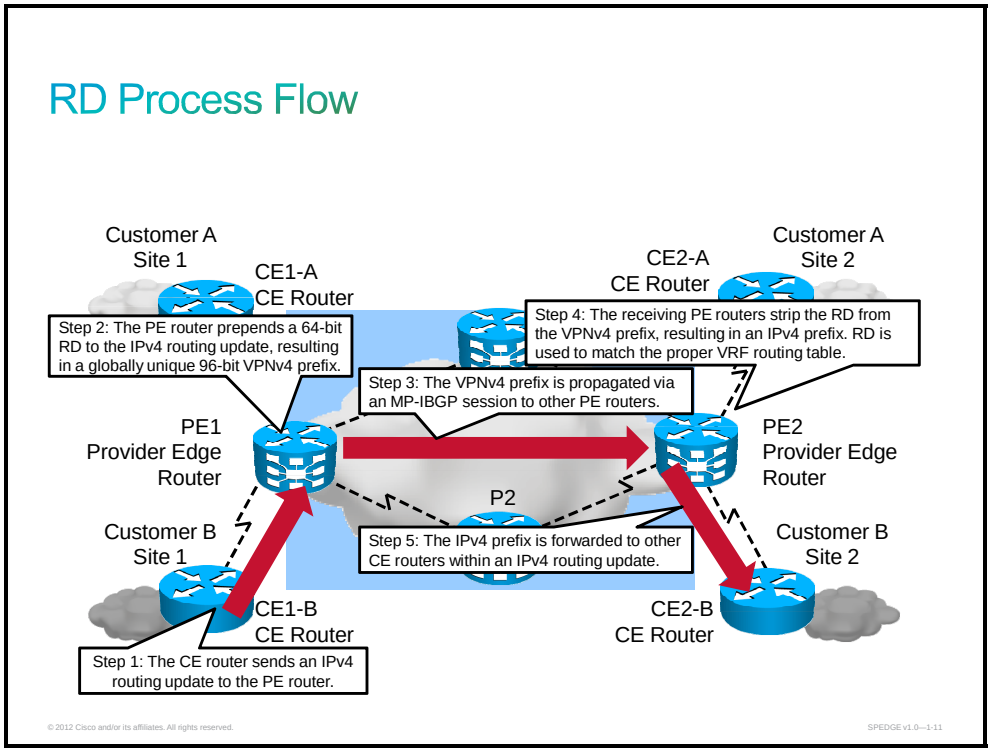
In the figure, the same IP prefix, 172.16.10.0/24, received from two different customers, is made unique by prepending different RD values, 1:100 and 1:101, before propagating the addresses as VPNv4 addresses on the PE router.

The protocol used for exchanging these VPNv4 routes between PE routers is MP-BGP; BGP that is capable of carrying VPNv4 (96-bit) prefixes in addition to other address families is called MP-BGP. The IGP requirement to implement Internal Border Gateway Protocol (IBGP) still holds in the case of an MPLS VPN implementation. Therefore, the PE router must run an IGP that provides Network Layer Reachability Information (NLRI) for IBGP if both PE routers are in the same AS.

MP-BGP is also responsible for the assignment of a VPN label. Packet forwarding in an MPLS VPN mandates that the router specified as the next hop in the incoming BGP update is the same router that assigns the VPN label.

An MP-BGP session between PE routers in a single BGP AS is called a Multiprotocol Internal Border Gateway Protocol (MP-IBGP) session and follows rules as in the implementation of IBGP with regards to BGP attributes. If the VPN extends beyond a single AS, VPNv4 routes will be exchanged between autonomous systems at the AS boundaries using a Multiprotocol External Border Gateway Protocol (MP-EBGP) session.

https://t.me/learningnets



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

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

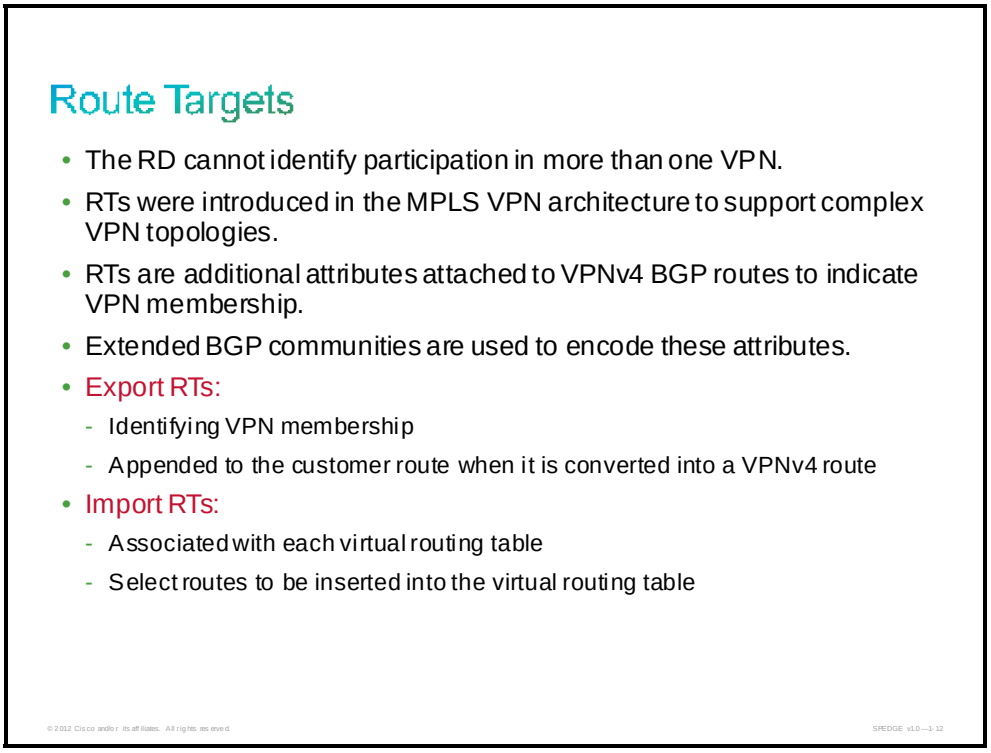
Step 2 The PE router prepends a 64-bit RD to the IPv4 routing update, resulting in a globally unique 96-bit VPNv4 prefix.

Step 3 The VPNv4 prefix is propagated via an MP-IBGP session to other PE routers.

Step 4 The receiving PE routers strip the RD from the VPNv4 prefix, resulting in an IPv4 prefix. RD is used to match the proper VRF routing table.

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

https://t.me/learningnets



Consider a scenario where some sites have to participate in more than one VPN. In such a scenario, the RDs cannot identify participation in more than one VPN.

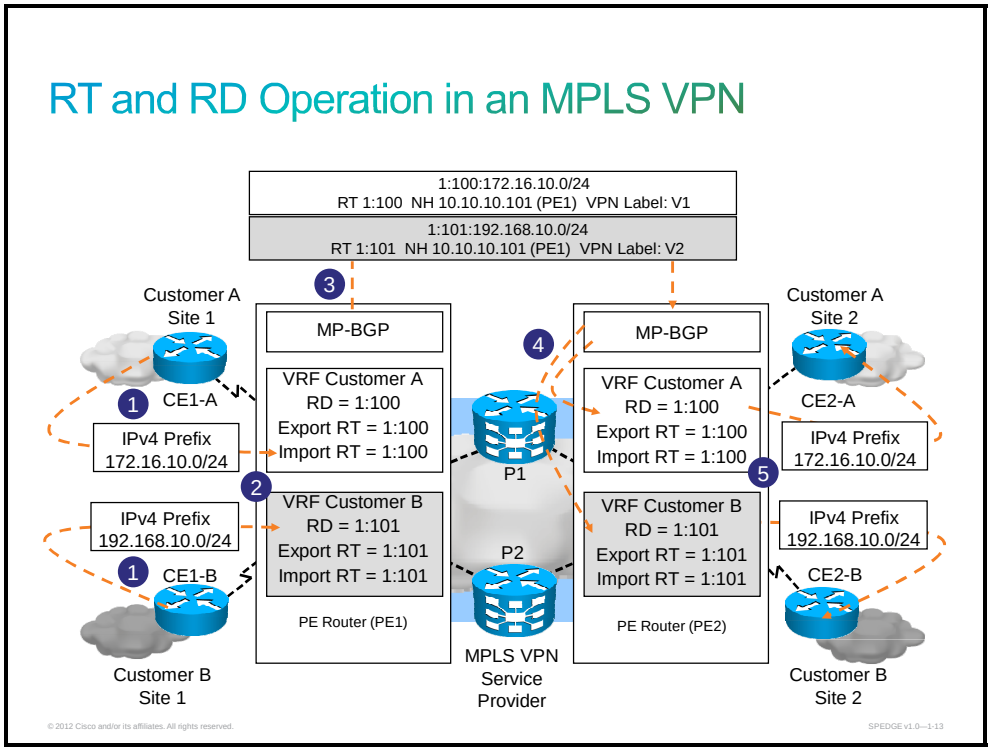
Route targets (RTs) were introduced into the MPLS VPN architecture to support identifying a site that participates in more than one VPN. RTs are additional identifiers used in the MPLS VPN that identify the VPN membership of the routes learned from that particular site. RTs are implemented by the use of extended BGP communities in which the higher-order 16 bits of the BGP extended community (64 total bits) are encoded with a value corresponding to the VPN membership of the specific site. When a VPN route learned from a CE router is injected into VPNv4 BGP, a list of VPN route target extended community attributes is associated with it.

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

When the VPNv4 routes are propagated to other PE routers, those routers need to select the routes to import into their virtual routing tables. This selection is based on import RTs. Each virtual routing table in a PE router can have a number of configured import RTs that identify the set of VPNs from which the virtual routing table is accepting routes.

When implementing complex VPN topologies, such as extranet VPN, Internet access VPNs, network management VPN, and so on, using MPLS VPN technology, the RT plays a pivotal role. A single prefix can be associated to more than one export route target when propagated across the MPLS VPN network. The RT can, as a result, be associated to sites that might be a member of more than one VPN.

https://t.me/learningnets



The following processes occur during route propagation in an MPLS VPN, as shown in the figure:

1. The prefix 172.16.10.0/24 is received from CE1-A, which is part of VRF Customer A on PE1. The prefix 192.168.10.0/24 is received from CE1-B, which is part of VRF Customer B on PE1.
2. For CE1-A, the PE1 associated an RD value of 1:100 and an export RT value of 1:100, as configured in the VRF definition on the PE router. For CE1-B, the PE1 associated an RD value of 1:101 and an export RT value of 1:101.
3. Routes learned from connected CE routers CE1-A are redistributed into the MP-BGP process on PE, where the prefix 172.16.10.0/24 is prepended with the RD value of 1:100 and appended with the route target extended community value (export RT) of 1:100 prior to sending the VPNv4 prefix as part of the MP-IBGP update between PE routers.

   Routes learned from connected CE routers CE1-B are redistributed into the MP-BGP process on PE1, where the prefix 192.168.10.0/24 is prepended with the RD value of 1:101 and appended with the route target extended community value (export RT) of 1:101 prior to sending the VPNv4 prefix as part of the MP-IBGP update between PE routers.

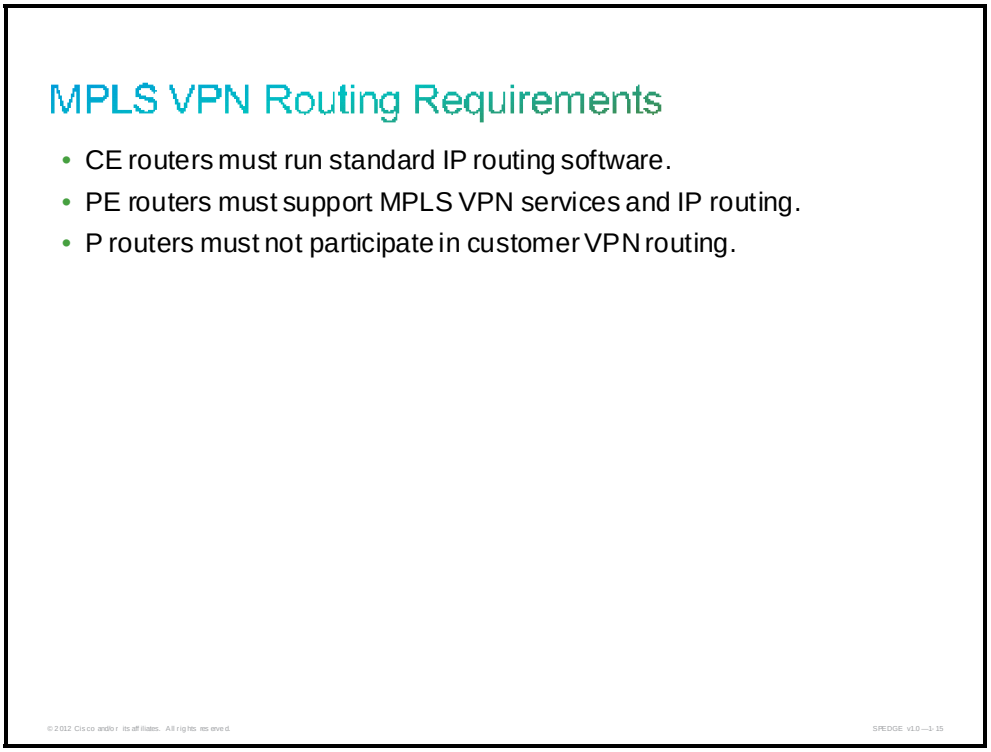
   The VPN label (4 bytes) is assigned for each prefix that is learned from the IGP process of the connected CE router within a VRF by the MP-BGP process of the PE router. MP-BGP running in the service provider MPLS domain thus carries the VPNv4 prefix (IPv4 prefix with prepended RD) in addition to the BGP route target extended community.

   The next hops on PE routers must not be advertised in the BGP process but must be learned from the IGP for MPLS VPN implementation. The VPN label is depicted by the entries V1 and V2 in the figure.
4. The MP-BGP update is received by the PE router PE2, and the route is stored in the appropriate VRF table for Customer A, based on the VPN label.
5. The received MP-BGP routes are redistributed into the VRF PE-CE routing processes, and the route is propagated to CE2-A.

https://t.me/learningnets

MPLS VPN Routing

This topic describes the end-to-end routing update flow for MPLS VPNs.

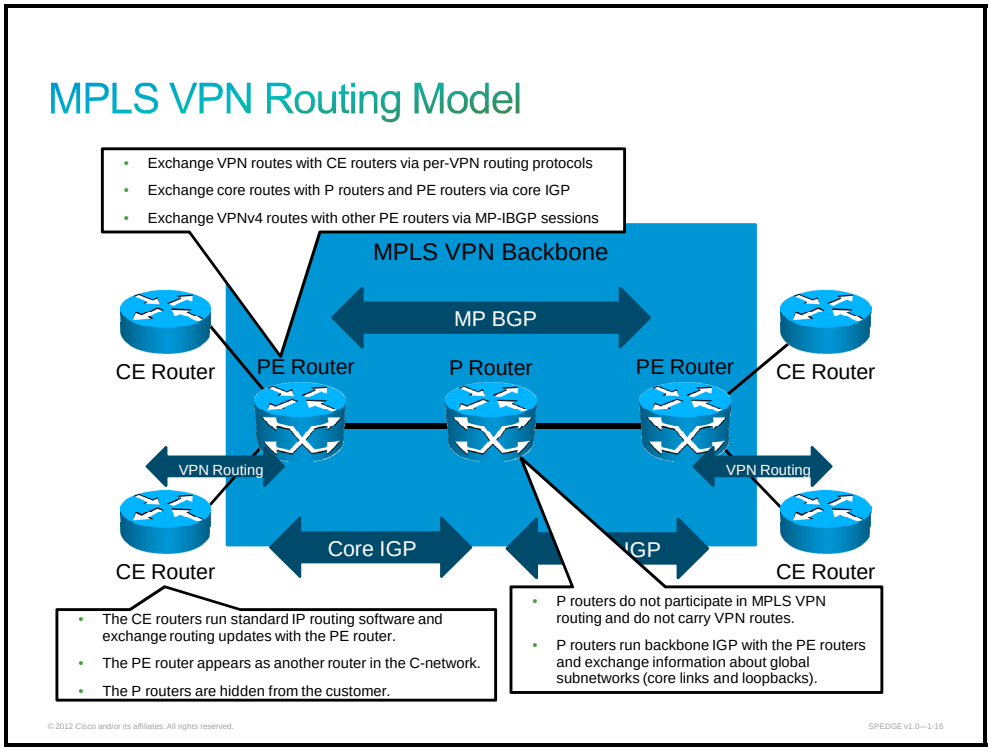


The designers of MPLS VPN technology were faced with these routing requirements:

- CE routers should not be MPLS VPN-aware; CE routers should run standard IP routing software.
- PE routers must support MPLS VPN services and traditional Internet services.

To make the MPLS VPN solution scalable, P routers must not participate in customer VPN routing. P routers use only label switching.

https://t.me/learningnets



The MPLS VPN backbone should look like a standard corporate backbone to the CE routers. The CE routers run standard IP routing software and exchange routing updates with the PE routers, which appear to them as normal routers in the customer network (C-network). An MPLS VPN implementation is very similar to a dedicated router peer-to-peer model implementation. From the perspective of a CE router, only IPv4 updates and data are forwarded to the PE router. The CE router does not need any specific configuration to enable it to be a part of an MPLS VPN domain. The only requirement on the CE router is a routing protocol (or a static default route) that enables the router to exchange IPv4 routing information with the connected PE router, which appears as normal router in the C-network.

From the customer perspective, the MPLS VPN backbone looks like an intracompany BGP backbone with PE routers performing route redistribution between individual sites and the core backbone. The standard design rules that are used for enterprise BGP backbones can be applied to the design of the C-network. The P routers are hidden from customer view; the internal topology of the BGP backbone is 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. The P routers run only a backbone IGP with other P routers and with PE routers, and exchange information about core subnetworks. BGP deployment on P routers is not needed for proper MPLS VPN operation; BGP deployment might be needed, however, to support traditional Internet connectivity that has not yet been migrated to MPLS.

The PE routers are the only routers in the MPLS VPN architecture that see all routing aspects of the MPLS VPN. PE routers can perform these exchanges:

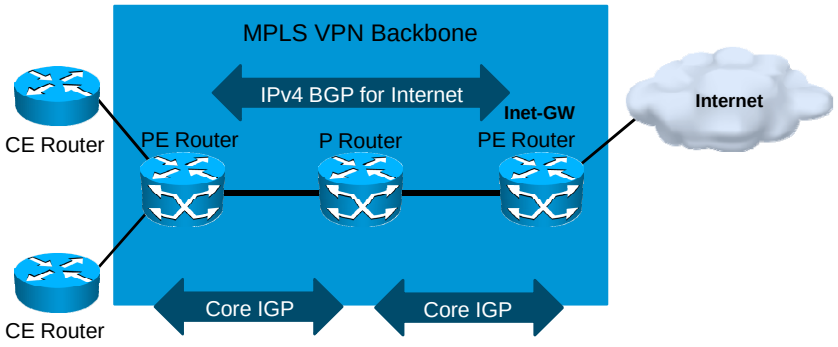
- PE routers exchange IPv4 VPN routes with CE routers via various routing protocols running in the virtual routing tables.
- PE routers exchange VPNv4 routes via MP-IBGP sessions with other PE routers.
- PE routers exchange core routes with P routers and other PE routers via core IGP.

https://t.me/learningnets

Support for Internet Routing

PE routers can run standard IPv4 BGP in the global routing table:

- PE routers exchange Internet routes with other PE routers.
- CE routers do not participate in Internet routing.
- P routers do not need to participate in Internet routing.



© 2012 Cisco and/or its affiliates. All rights reserved.

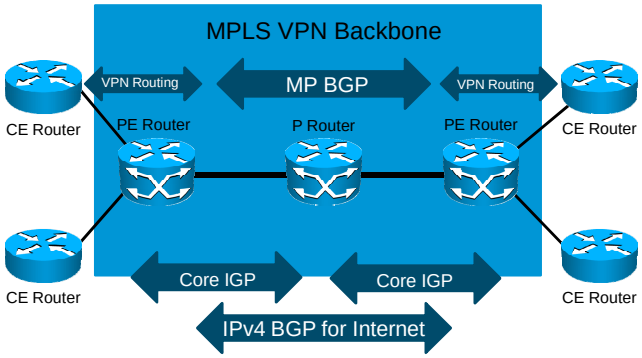
The routing requirements for PE routers also extend to supporting Internet connectivity—PE routers need to exchange Internet routes with other PE routers. The CE routers cannot participate in Internet routing if the Internet routing is performed in global address space. The P routers could participate in Internet routing; however, Internet routing should be disabled on the P routers to make the network core more stable.

https://t.me/learningnets

Routing Tables on PE Routers

PE routers contain a number of routing tables:

- The global routing table contains core routes (filled with core IGP) and Internet routes (filled with IPv4 BGP).
- The VRF tables contain routes for sites of identical routing requirements from local (IPv4 VPN) and remote (VPNv4 via MP-BGP) CE routers.

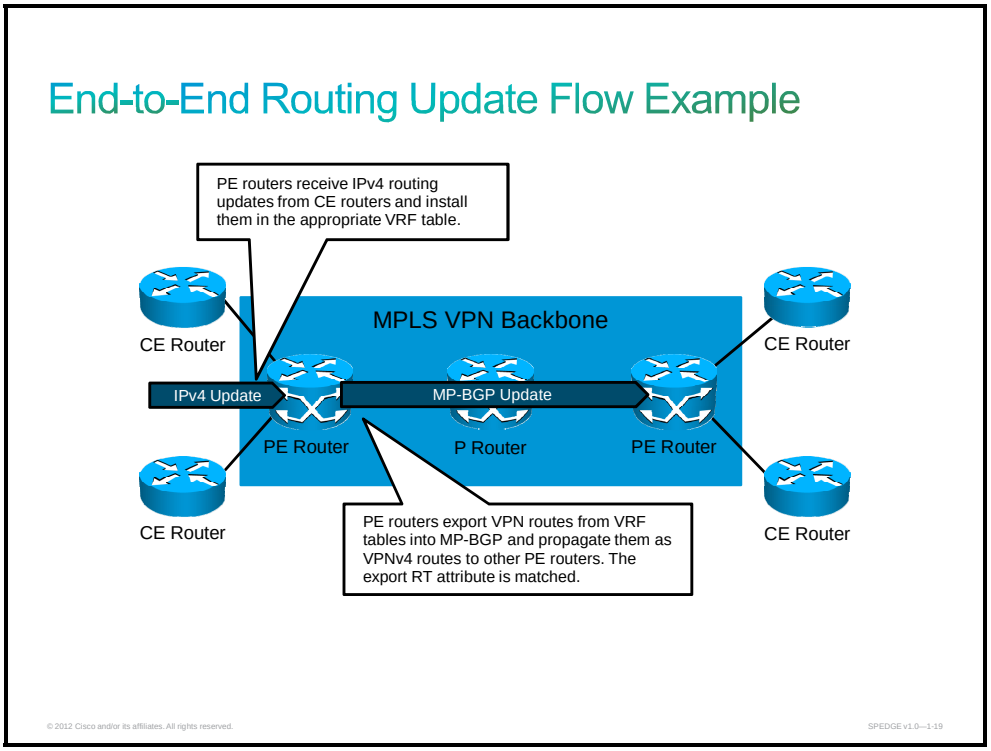


© 2012 Cisco and/or its affiliates. All rights reserved.

The PE routers fulfill various routing requirements imposed on them by using a number of IP routing tables. Here are some examples:

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

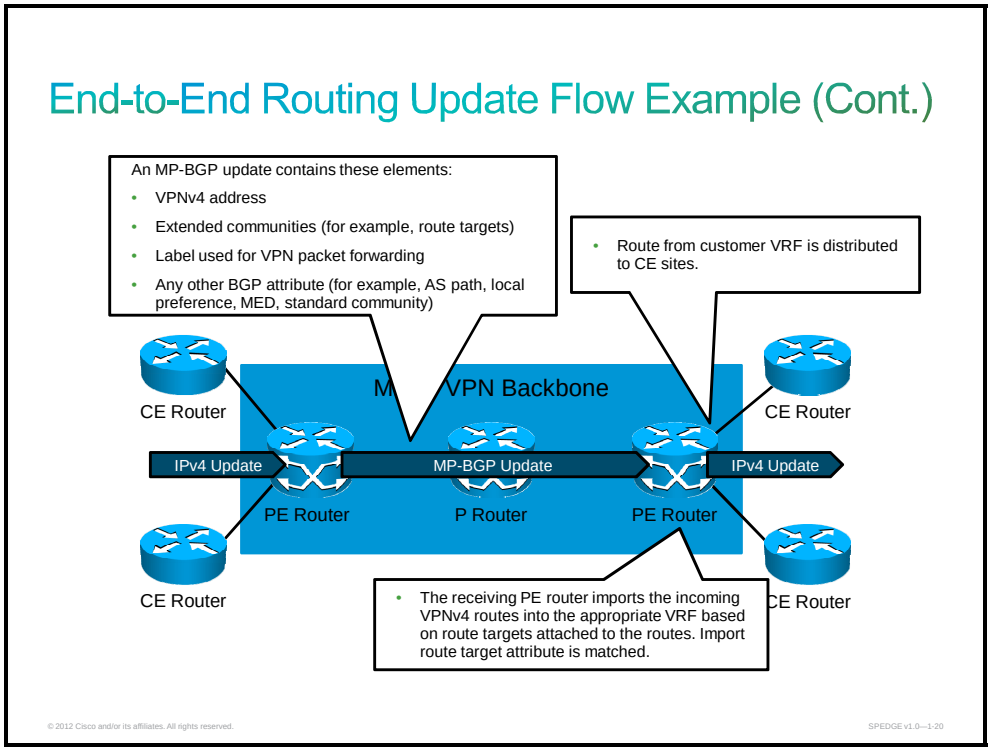
https://t.me/learningnets



The figure shows how PE routers receive IPv4 routing updates from the CE routers and install the updates in the appropriate VRF table. Cisco IOS, IOS XE, and IOS XR Software currently supports RIP version 2 (RIPv2) (multiple contexts), EIGRP (multiple contexts), OSPF version 2 (OSPFv2) (multiple processes), Intermediate System-to-Intermediate System (IS-IS) (multiple contexts), and BGP4 (multiple contexts) as routing protocols that can be used per-VRF to exchange customer routing information between CE routers and PE routers. The VRF interfaces on PE routers can be either logical or physical, but each interface can be assigned to only one VRF.

The customer routes from VRF tables are exported as VPNv4 routes into MP-BGP and propagated to other PE routers. The MP-BGP sessions between the PE routers are IBGP sessions and are subject to the IBGP split-horizon rules. Either a full mesh of MP-IBGP sessions is required between PE routers, or route reflectors need to be used.

https://t.me/learningnets



An MP-BGP update exchange between PE routers contains these elements:

- VPNv4 address
- Extended BGP communities (for example, RTs are required)
- Label used for VPN packet forwarding
- Mandatory BGP attributes (for example, AS path). Optionally, the MP-BGP update can contain any other BGP attribute, such as local preference, multi-exit discriminator (MED), or standard BGP community.

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

The RTs that are attached to a route and the import RTs that are configured in the VRF direct the propagation of the routes to the CE router. Incoming VPNv4 routes are imported into VRFs on the receiving PE router only if at least one RT attached to the route matches at least one import RT that is configured in the VRF.

https://t.me/learningnets

MPLS VPN Forwarding Mechanisms

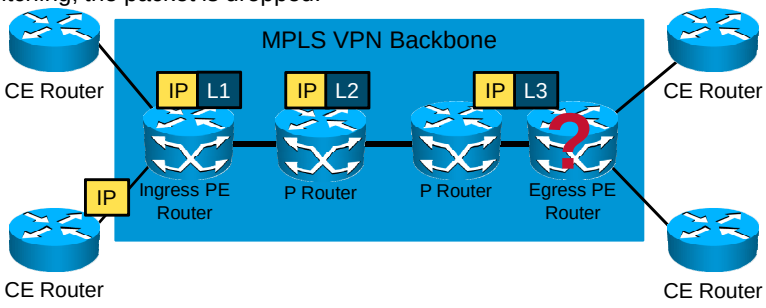
This topic describes VPN label propagation between PE routers and the MPLS VPN end-to-end forwarding mechanism.

VPN Packet Forwarding Across an MPLS VPN Backbone—Approach 1

Approach 1: The PE routers will label the VPN packets with an LDP label for the egress PE router, and forward the labeled packets across the MPLS backbone.

Results:

- The P routers perform the label switching, and the packet reaches the egress PE router.
- Because the egress PE router does not know which VRF to use for packet switching, the packet is dropped.



© 2012 Cisco and/or its affiliates. All rights reserved.

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

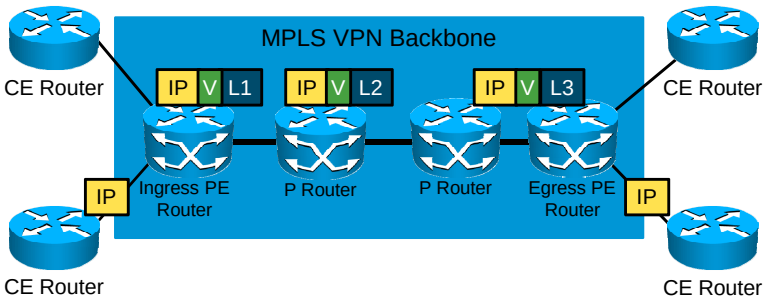
https://t.me/learningnets

VPN Packet Forwarding Across an MPLS VPN Backbone—Approach 2

Approach 2: The PE routers will label the VPN packets with a label stack, using the LDP label for the egress PE router as the top label, and the VPN label assigned by the egress PE router as the second label in the stack.

Results:

- The P routers perform label switching using the top label, and the packet reaches the egress PE router. The top label is removed.
- The egress PE router performs a lookup on the VPN label and forwards the packet toward the CE router.



© 2012 Cisco and/or its affiliates. All rights reserved.

An MPLS label stack can be used to tell the egress PE router what to do with the VPN packet. When using the label stack, the ingress PE router labels the incoming IP packet with two labels. The top label in the stack is the LDP label for the egress PE router; this label guarantees that the packet will traverse the MPLS VPN backbone and arrive at the egress PE router. The second label in the stack is assigned by the egress PE router, and tells how to forward the incoming VPN packet. The second label could point directly toward an outgoing interface, in which case the egress PE router would perform label lookup only on the VPN packet. The second label could also point to a VRF, in which case the egress PE router would first perform a label lookup to find the target VRF, and then perform an IP lookup within the VRF.

Both methods of implementing second labels are used in Cisco IOS, IOS XE, and IOS XR Software. The second label in the stack points toward an outgoing interface whenever the CE router is the next hop of the VPN route. The second label in the stack points to the VRF table for aggregate VPN routes, VPN routes pointing to a null interface, and routes for directly connected VPN interfaces.

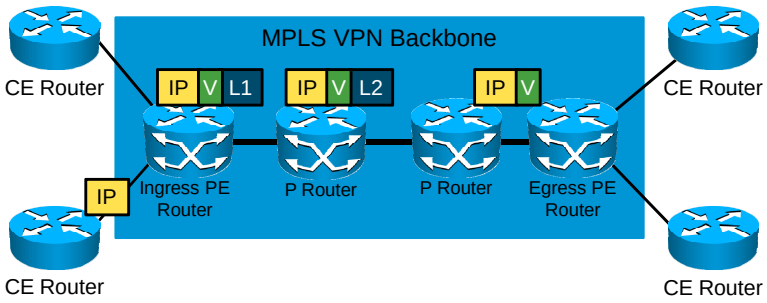
The two-level MPLS label stack satisfies these MPLS VPN forwarding requirements:

- The P routers perform label switching on the LDP-assigned label toward the egress PE router.
- The egress PE router performs label switching on the second label (which it has previously assigned) and either forwards the IP packet toward the CE router or performs another IP lookup in the VRF that is pointed to by the second label in the stack.

https://t.me/learningnets

VPN PHP

- PHP on the LDP label can be performed on the last P router.
- The egress PE router performs label lookup only on the VPN label, resulting in faster and simpler label lookup.
- IP lookup is performed only once—in the ingress PE router.



© 2012 Cisco and/or its affiliates. All rights reserved.

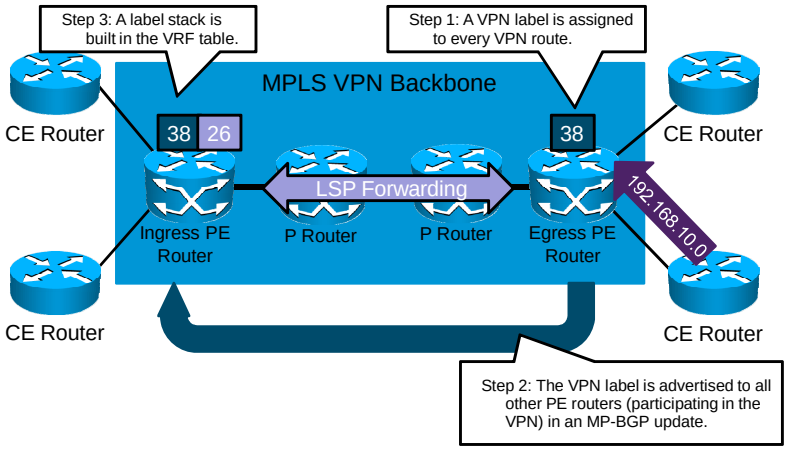
Penultimate hop popping (PHP) is the removal of the top label in the stack on the hop prior to the egress router. PHP can be performed in frame-based MPLS networks. In these networks, the last P router in the label-switched path (LSP) tunnel pops the LDP label (as previously requested by the egress PE router through LDP), and the PE router receives a labeled packet that contains only the VPN label. In most cases, a single label lookup performed on that packet in the egress PE router is enough to forward the packet toward the CE router. The full IP lookup through the forwarding information base (FIB) is performed only once, in the ingress PE router, even without PHP.

https://t.me/learningnets

VPN Label Propagation

Question: How will the ingress PE router get the second label in the label stack from the egress PE router?

Answer: Labels are propagated in MP-BGP VPNv4 routing updates.



© 2012 Cisco and/or its affiliates. All rights reserved.

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

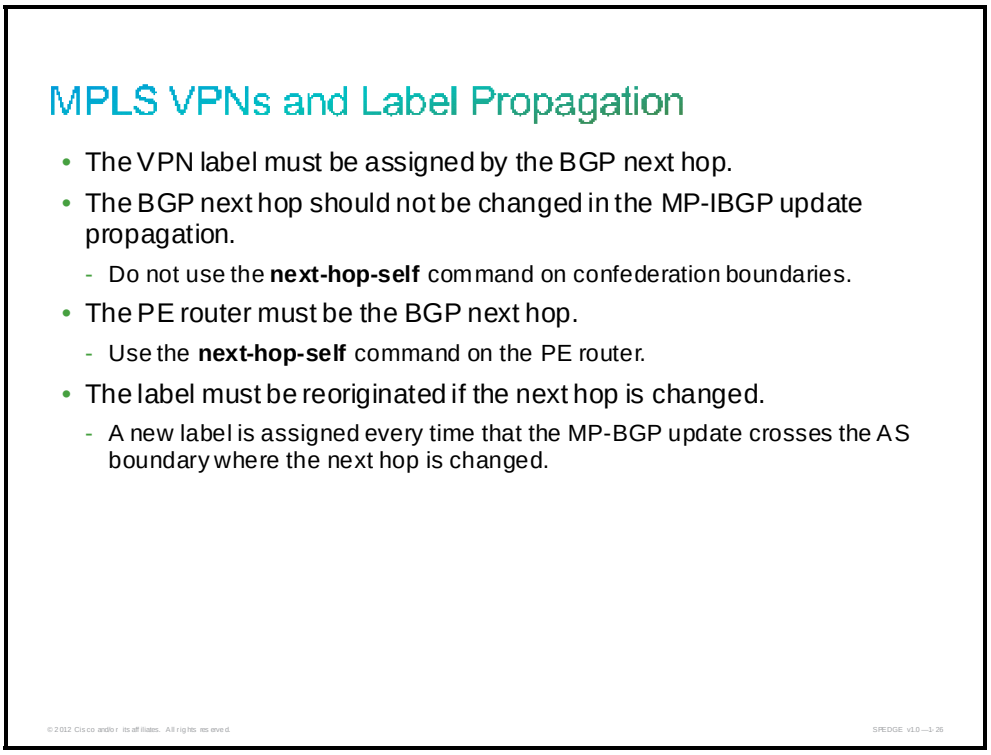
These steps describe the label propagation between PE routers:

Step 1 The egress PE router assigns a label to every VPN route received from the attached CE routers and to every summary route summarized inside the PE router. This label is then used as the second label in the MPLS label stack by the ingress PE routers when labeling VPN packets. In the figure, the VPN label 38 for destination 192.168.10.0 is assigned by the egress PE router. The VPN labels that are assigned locally by the PE router can be inspected with the Cisco IOS, IOS XE, and IOS XR **show mpls forwarding vrf** *vrf-name* command.

Step 2 The VPN labels that are assigned by the egress PE routers are advertised to all other PE routers (participating in the VPN) together with the VPNv4 prefix in MP-BGP updates. The labels can be inspected with the Cisco IOS and IOS XE **show ip bgp vpnv4 all labels** command or the Cisco IOS XR **show bgp vpnv4 unicast labels** command on the ingress PE router. The routes that have an input label but no output label are the routes received from the CE routers (and the input label was assigned by the local PE router). The routes with an output label but no input label are the routes received from the other PE routers (and the output label was assigned by the remote PE router).

https://t.me/learningnets

Step 3 The ingress PE router has two labels associated with a remote VPN route: a label for the next hop assigned by the next-hop P router via LDP—and taken from the local label information base (LIB)—and also the label assigned by the remote PE router and propagated via the MP-BGP update. Both labels are combined in a label stack and installed in the VRF table. The label stack in the VRF table can be inspected using the Cisco IOS and IOS XE **show ip cef vrf** *vrf-name* **detail** command or the Cisco IOS XR **show cef vrf** *vrf-name* **detail** command. The *tags imposed* values in the output display the MPLS label stack. The first label in the MPLS label stack is the LDP label forwarded toward the egress PE router, and the second label is the VPN label advertised by the egress PE router.



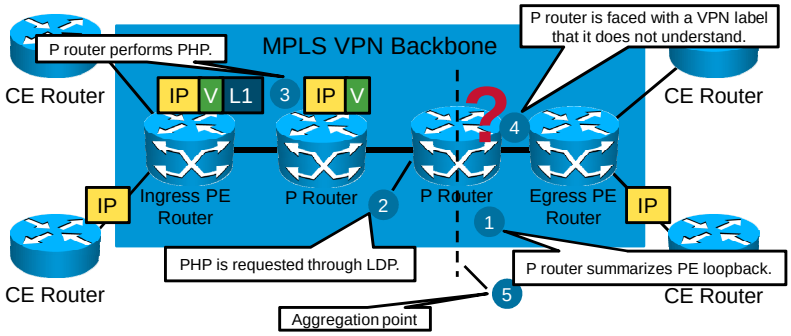
MPLS VPN packet forwarding works correctly only if the router specified as the BGP next hop in the incoming BGP update is the same PE router that assigned the second label in the label stack. Here are three scenarios that can cause the BGP next hop to be different from the IP address of the PE router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

- If the customer route is received from the CE router via an External Border Gateway Protocol (EBGP) session, the next hop of the VPNv4 route is still the IP address of the CE router (the BGP next hop of an outgoing IBGP update is always identical to the BGP next hop of the incoming EBGP update). You must configure the **next-hop-self** command on the MP-BGP sessions between PE routers to make sure that the BGP next hop of the VPNv4 route is always the IP address of the PE router, regardless of the routing protocol used between the PE router and the CE router.
- The BGP next hop should not change inside an AS.
- The BGP next hop is always changed on an EBGP session. If the MPLS VPN network spans multiple public autonomous systems, special provisions must be made in the AS boundary routers to reoriginate the VPN label at the same time that the BGP next hop is changed.

https://t.me/learningnets

MPLS VPNs and Packet Forwarding

- The VPN label of the BGP route is understood by only the egress PE router.
- An end-to-end LSP tunnel is required between the ingress and egress PE routers.
- BGP next-hop addresses must be IGP routes.
  - LDP labels will be assigned to addresses in the global routing table.
  - LDP labels are not assigned to BGP routes (BGP routes receive VPN labels).
- BGP next hops announced in IGP must not be summarized in the core network.
  - Summarization breaks the LSP tunnel.



© 2012 Cisco and/or its affiliates. All rights reserved.

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

Here are two scenarios that could cause the LSP tunnel between PE routers to break:

- If the IP address of the PE router is announced as a BGP route, it will have no corresponding LDP label, and the label stack will not be built correctly. The IP address of the PE router must be announced in the global routing table.
- If the P routers perform summarization of the address range within which the IP address of the egress PE router lies, the LSP tunnel will be disrupted at the summarization point.

In the figure, the P router summarizes the loopback address of the egress PE router. The LSP tunnel is broken at a summarization point, so the summarizing router needs to perform full IP lookup. In an MPLS network, the P router would request PHP for the summary route, and the upstream P router (or a PE router) would remove the LDP label, exposing the VPN label to the P router. Because the VPN label is assigned not by the P router but by the egress PE router, the label will not be understood by the P router and the VPN packet will be dropped or misrouted.

https://t.me/learningnets

Summary

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

Summary

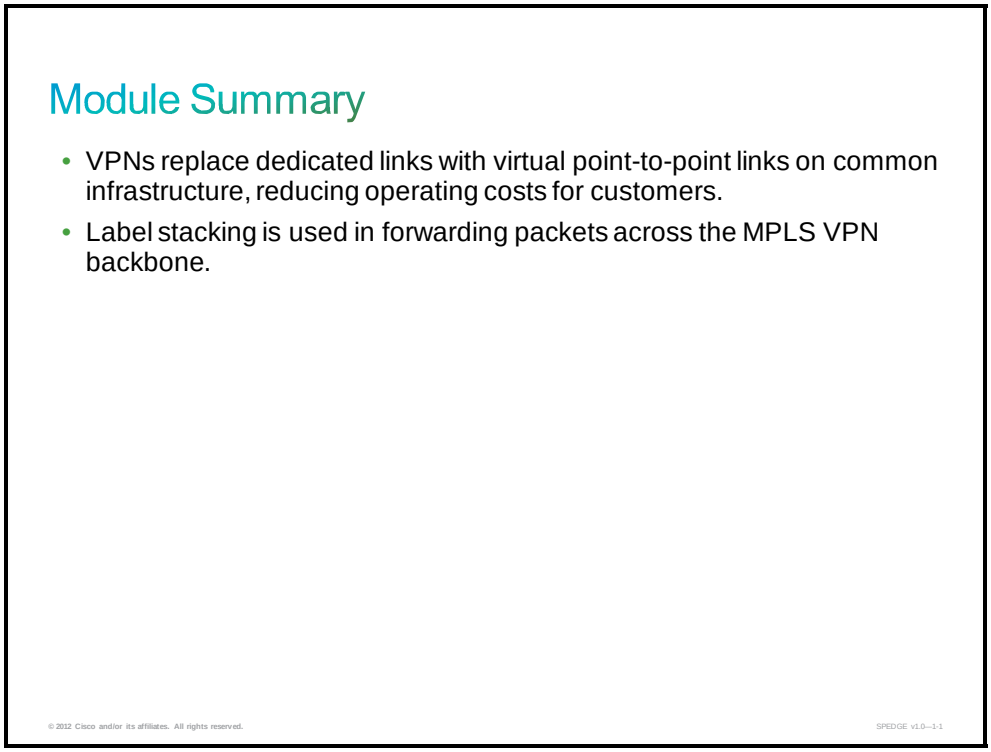
- The most scalable method of exchanging customer routes across a provider network is the use of an MP-BGP between PE routers.
  - RDs transform non-unique 32-bit addresses into 96-bit unique addresses.
  - RTs are used to identify VPN membership in overlapping topologies.
- In MPLS VPNs:
  - CE routers run standard routing protocols to the PE routers.
  - PE routers provide the VPN routing and services via MP-BGP.
  - P routers do not participate in VPN routing, and only provide core IGP backbone routing to the PE routers.
- PE routers forward packets across the MPLS VPN backbone using label stacking.

© 2012 Cisco and/or its affiliates. All rights reserved.

https://t.me/learningnets

Module Summary

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



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

https://t.me/learningnets

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Module Self-Check

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

Q1) Which three options are advantages of VPNs? (Choose three.) (Source: Introducing VPNs)

A) cost savings
B) scalability
C) improved security
D) complexity
E) simplified routing

Q2) Which two options are Layer 2 overlay VPN technologies? (Choose two.) (Source: Introducing VPNs)

A) Frame Relay
B) GRE
C) SSL VPN
D) ATM

Q3) Which two options are Layer 3 overlay VPN technologies? (Choose two.) (Source: Introducing VPNs)

A) Frame Relay
B) DMVPN
C) IPsec
D) ATM

Q4) Which of these is a peer-to-peer VPN technology? (Source: Introducing VPNs)

A) DMVPN
B) GET VPN
C) IPsec
D) L2TPv3

Q5) Which routers are MPLS VPNs aware of? (Source: Introducing MPLS VPNs)

Q6) Which protocol is used to transport customer routes directly between PE routers? (Source: Introducing MPLS VPNs)

A) RIP
B) VPN
C) BGP
D) OSPF

Q7) In which two ways do MPLS VPNs support overlapping customer address spaces? (Choose two.) (Source Introducing MPLS VPNs)

A) by implementing unique RDs for each customer
B) by implementing unique RTs for each customer
C) by implementing different LSPs for each customer
D) by implementing virtual routing spaces for each customer

https://t.me/learningnets

Q8) Why do MPLS VPNs implement route targets? (Source: Introducing MPLS VPNs)

A) to identify different customer VPNs
B) to allow a site to participate in more than one VPN
C) to convert a customer address to an MP-BGP address
D) to convert a non-unique IP address into a unique VPNv4 address

Q9) Which type of routers exchange VPNv4 routes? (Source: Introducing MPLS VPNs)

A) P
B) CE
C) PE
D) core

Q10) Which protocol would a PE router use to support an existing Internet routing scheme? (Source: Introducing MPLS VPNs)

A) IS-IS
B) EIGRP
C) BGP4
D) BGP VPNv4

Q11) What is the effect of an MPLS VPN on CE routers? (Source: Introducing MPLS VPNs)

A) The CE routers must support BGP.
B) The CE routers must run a link-state protocol.
C) The CE routers can run any standard IP routing protocol.
D) The IGP of the CE routers must be upgraded to a VPN-aware IGP.

Q12) Why would IPv4 routing be enabled on the PE router? (Source: Introducing MPLS VPNs)

A) to support the MPLS VPN route update
B) to support the MPLS VPN route target exports
C) to support an existing Internet routing scheme
D) to support the transport of MP-BGP extended communities

Q13) Which two types of routes would an MPLS VPN install into the VRF? (Choose two.) (Source: Introducing MPLS VPNs)

A) those routes received via an IPv4 update
B) those routes received via a VPNv4 update
C) those routes received via the core IGP update
D) those routes received via the customer IGP update

Q14) Which protocol is used to transport VPN labels between PE routers? (Source: Introducing MPLS VPNs)

A) LDP
B) RSVP
C) MP-BGP
D) the core IGP

Q15) How can P routers forward VPN packets if they do not have VPN routes? (Source: Introducing MPLS VPNs)

A) They forward based on the LSP label.
B) They forward based on the VPN label.
C) They forward based on the MP-BGP next hop.
D) They forward based on a routing table lookup of the IP address.

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Q16) Which router assigns the VPN label? (Source: Introducing MPLS VPNs)

A) P
B) egress CE
C) egress PE
D) ingress CE
E) ingress PE

Q17) What is used to identify the label that will be used to transport the VPN packet to the egress router? (Source: Introducing MPLS VPNs)

A) the IGP least-cost path
B) the EBGP next-hop address
C) the MP-IBGP next-hop address
D) the VPN label entry in the LFIB

https://t.me/learningnets

Module Self-Check Answer Key

Q1) A, B, C

Q2) A, D

Q3) B, C

Q4) B

Q5) PE routers

Q6) C

Q7) A, D

Q8) B

Q9) C

Q10) C

Q11) C

Q12) C

Q13) B, D

Q14) C

Q15) A

Q16) C

Q17) C

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Module 2

MPLS Layer 3 VPNs

Overview

This module covers Multiprotocol Label Switching (MPLS) VPN implementation on Cisco IOS Software platforms. The module describes the concepts of virtual routing and forwarding (VRF) tables, the interaction between customer-to-provider routing protocols, and Multiprotocol Border Gateway Protocol (MP-BGP) in the service provider backbone. It also describes advanced MPLS VPN-specific routing protocol features. The module continues with a description of the MPLS VPN monitoring and debugging commands that are available on Cisco IOS platforms and describes troubleshooting, including failure scenarios, symptoms, and remedial action. At the end, this module describes IPv6 service provider deployment strategies.

Module Objectives

Upon completing this module, you will be able to configure, monitor, and troubleshoot VPN operations and identify IPv6 strategies in service provider environments. You will be able to meet these objectives:

- Configure VRF tables and MP-BGP sessions between PE routers
- Configure small-scale routing protocols (static, RIP, and EIGRP) between CE and PE routers
- Configure OSPF and BGP as the routing protocol between CE and PE routers and explain how to troubleshoot MPLS VPN operations
- Describe various methods that are used to deploy IPv6 over MPLS

https://t.me/learningnets

https://t.me/learningnets

Lesson 1

Implementing MPLS Layer 3 VPN Backbones

Overview

This lesson introduces the virtual routing and forwarding (VRF) table, the major data structure that is associated with Multiprotocol Label Switching (MPLS) VPN implementation on Cisco IOS, IOS XE, and IOS XR platforms. The lesson describes the MPLS VPN attributes that are associated with a VRF instance and explains the need for routing protocol contexts and the interaction of routing protocol contexts, VRFs, and Multiprotocol Border Gateway Protocol (MP-BGP).

Having a clear understanding of how information is exchanged using VRFs and routing protocol contexts will make it easier to configure VRFs in your network.

This lesson also explains how to configure VRF tables, listing the configuration tasks, syntax, and definitions of commands that are used to create VRFs. It provides an example of a VPN configuration.

It is important to know how to configure and apply a VRF table onto a routing interface. It is essential to understand the command syntax for the configurations that you want to deploy in your network. This lesson provides you with the information that will enable you to succeed at such tasks.

Objectives

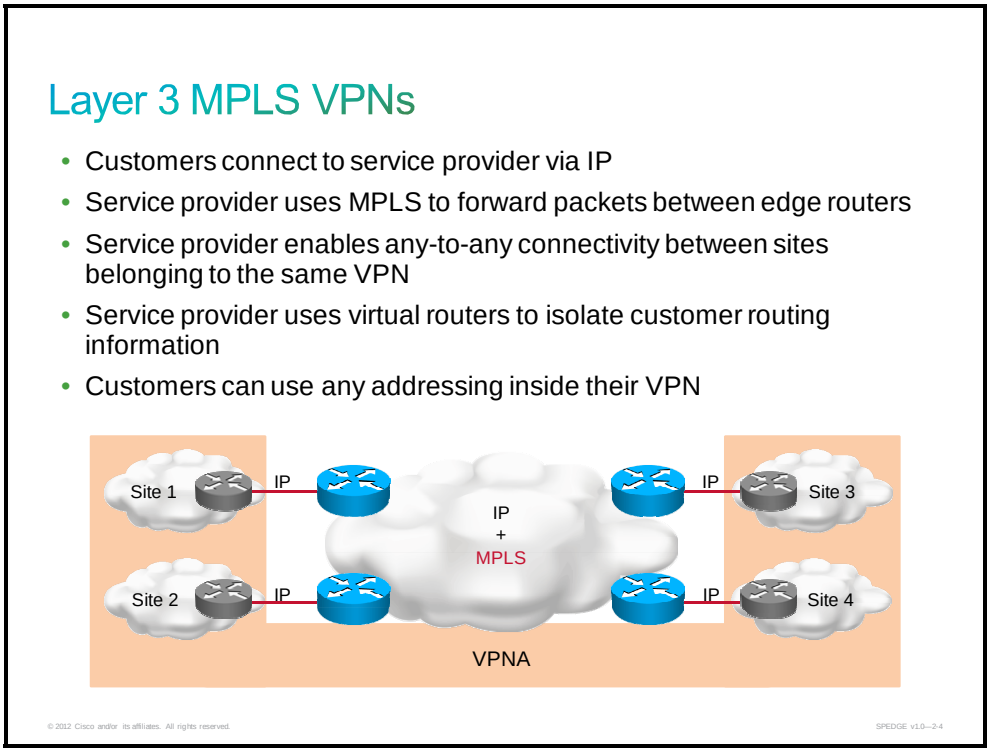
Upon completing this lesson, you will be able to describe the VRF table and other MPLS VPN attributes that are associated with a VRF instance. You will be able to meet these objectives:

- Describe VRF
- Enable VRF
- Enable MP-BGP

https://t.me/learningnets

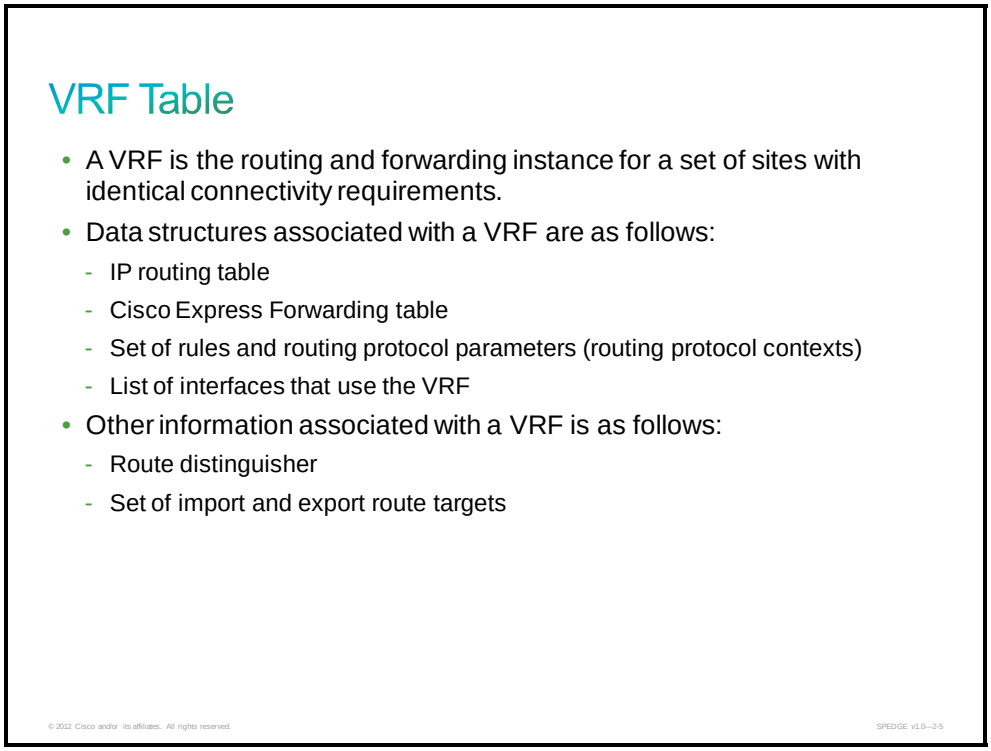
Virtual Routing and Forwarding

This topic describes the VRF table and describes the other MPLS VPN attributes that are associated with a VRF instance.



The main characteristic of Layer 3 MPLS VPNs is that customers connect to the service provider via IP. They need to establish IP routing (static or dynamic) to exchange routing information between customer sites that belong to the same VPN. Because different customers might use the same private IP address ranges, the service provider cannot perform normal IP forwarding. Instead, service providers must use MPLS to ensure isolation in the data plane between packets belonging to different customers but potentially having the same IP addresses. Virtual routers (VRF instances) are used on service provider routers to isolate customer routing information. MPLS seamlessly provides any-to-any connectivity between sites that belong to the same VPN.

https://t.me/learningnets



The major data structure that is associated with MPLS VPN implementation on Cisco IOS platforms is the VRF table. This data structure encompasses an IP routing table that is identical in function to the following:

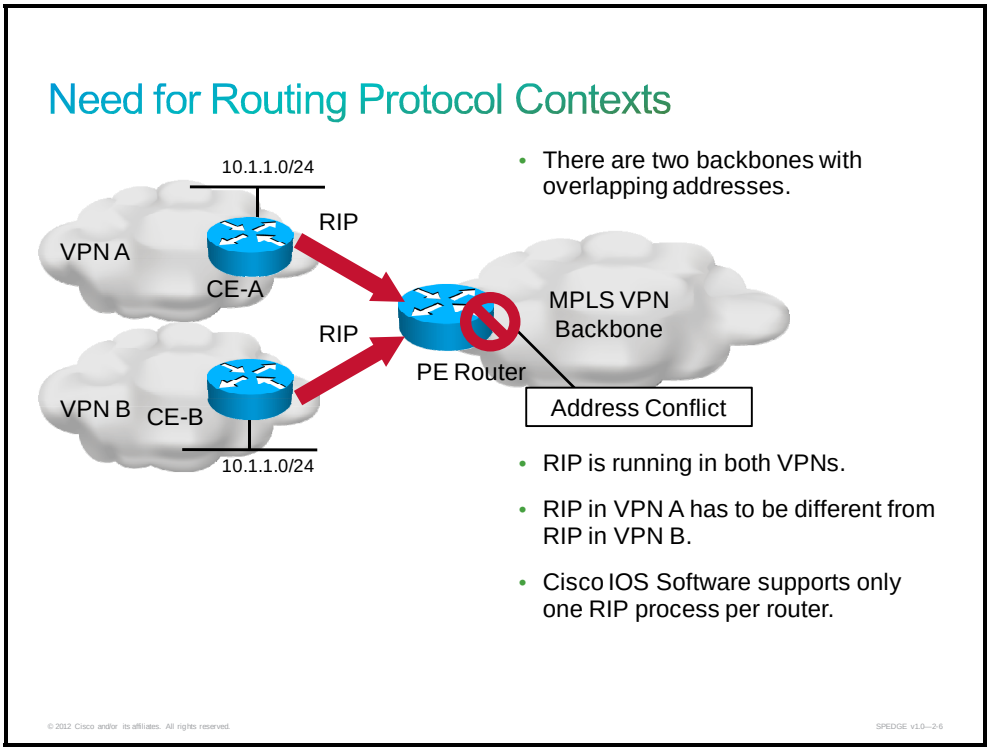
- The global IP routing table in Cisco IOS Software
- A Cisco Express Forwarding table that is identical in function to the global Cisco Express Forwarding table (forwarding information base [FIB])
- Specifications for routing protocols running inside the VRF instance

A VRF is a routing and forwarding instance that you can use for a single VPN site or for many sites connected to the same provider edge (PE) router, if and only if these sites share exactly the same connectivity requirements.

Other MPLS VPN attributes that are associated with a VRF table are as follows:

- The route distinguisher (RD), which is prepended (for example, RD + IP address) to all routes that are exported from the VRF into the global VPNv4 (also called VPN IPv4) BGP table
- A set of export route targets (RTs), which are attached to any route that is exported from the VRF
- A set of import RTs, which are used to select VPNv4 routes that are to be imported into the VRF

https://t.me/learningnets



Traditional Cisco IOS Software can support a number of different routing protocols. In some cases, even several completely isolated copies of the same routing protocol are supported. For example, several Open Shortest Path First (OSPF) processes can be used.

It is important to understand that for several important routing protocols, such as Routing Information Protocol (RIP), Enhanced Interior Gateway Routing Protocol (EIGRP), Intermediate System-to-Intermediate System (IS-IS), and Border Gateway Protocol (BGP), Cisco IOS Software supports only a single copy of the protocol running in the router. These protocols cannot be used directly between PE and customer edge (CE) routers in VPN environments because each VPN (or, more precisely, each VRF) needs a separate, isolated copy of the routing protocol to prevent undesired route leakage between VPNs. Furthermore, VPNs can use overlapping IP address spaces (for example, each VPN could use subnetworks of network 10.1.1.0/24), which would also lead to routing confusion if all VPNs shared the same copy of the routing protocol.

https://t.me/learningnets

VPN-Aware Routing Protocols

Routing context = routing protocol run in one VRF

- Supported by VPN-aware routing protocols: External BGP (EBGP), EIGRP, OSPF, RIPv2, IS-IS, static routes
- Implemented as several instances of a single routing process (EIGRP, EBGP, RIPv2, IS-IS) or as several routing processes (OSPF)
- Independent per-instance router variables for each instance

© 2012 Cisco and/or its affiliates. All rights reserved.

"Routing contexts" were introduced in Cisco IOS Software to support the need for separate isolated copies of VPN routing protocols. Routing contexts can be implemented as separate routing processes (in OSPF), similar to a traditional Cisco IOS Software implementation, or as separate isolated instances of the same routing protocol.

If routing contexts are implemented as instances of the same routing protocol, each instance contains its own independent routing protocol parameters. Examples would include networks over which the routing protocol is run, timers, authentication parameters, passive interfaces, and neighbors. This independence allows the network designer to have maximum flexibility in implementing routing protocols between PE and CE routers.

https://t.me/learningnets

VRF Table

- Contains routes that should be available to a particular set of sites
- Analogous to standard Cisco IOS Software routing table; supports same set of mechanisms
- VPN interfaces (physical interface, subinterfaces, logical interfaces) assigned to VRFs:
  - Many interfaces per VRF
  - Each interface assignable to only one VRF

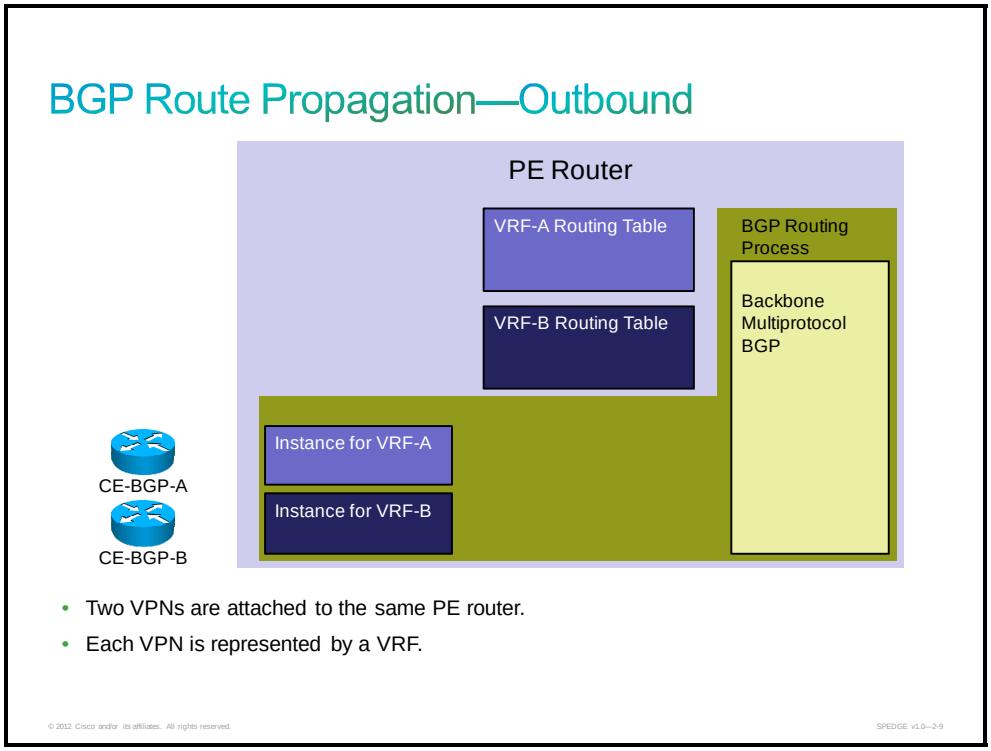
© 2012 Cisco and/or its affiliates. All rights reserved.

The routes that are received from VRF routing protocol instances or from dedicated VRF routing processes are inserted into the IP routing table that is contained within the VRF. This IP routing table supports exactly the same set of mechanisms as the standard Cisco IOS Software routing table. These mechanisms include filter mechanisms (distribute lists or prefix lists) and interprotocol route-selection mechanisms (administrative distances).

The per-VRF forwarding information base (FIB) table is built from the per-VRF routing table. This table is used to forward all the packets that are received through the interfaces that are associated with the VRF. Any interface can be associated with a VRF, whether it is a physical interface, subinterface, or logical interface, as long as it supports Cisco Express Forwarding switching.

There is no limit to the number of interfaces that can be associated with one VRF (other than the number of interfaces that are supported by the router). However, each interface can be assigned to only one VRF because the router needs to uniquely identify the forwarding table to be used for packets that are received over an interface.

https://t.me/learningnets

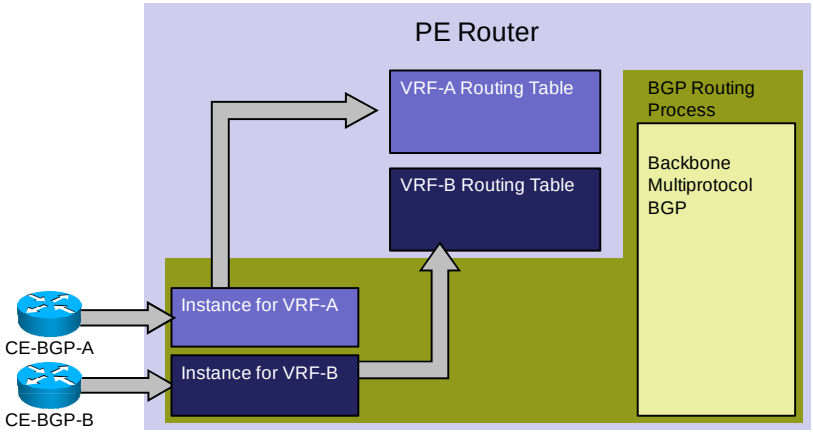


This figure and the next figures illustrate the interactions between VRF instances of routing processes, VRF routing tables, and the global VPNv4 BGP routing process.

The network contains two VPN customers. Ordinarily, the customer sites would be connected to a number of PE routers. This example focuses on only a single PE router, which contains two VRFs—one for each customer. Each customer is connected to the PE router, which is running BGP. CE-BGP-A is the CE router for customer A and is associated with VRF-A (VPN-A). CE-BGP-B is the CE router for customer B and is associated with VRF-B (VPN-B). Both CE routers are using BGP for exchanging routes with the PE router.

https://t.me/learningnets

BGP Route Propagation—Outbound (Cont.)



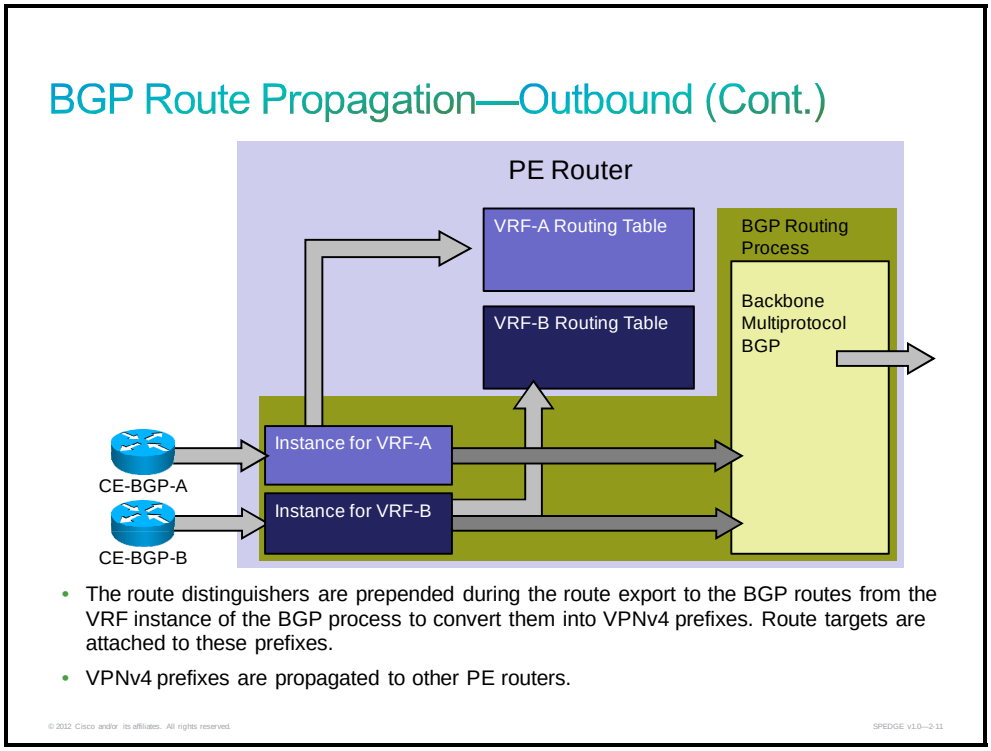
- BGP-speaking CE routers announce their prefixes to the PE router via BGP.
- The instance of the BGP process associated with the VRF of the PE-CE interface collects the routes and inserts them into the VRF routing table.

© 2012 Cisco and/or its affiliates. All rights reserved.

The CE routers are BGP neighbors of the PE router. The BGP-speaking CE routers announce their networks via External Border Gateway Protocol (EBGP) sessions to the PE router. The PE router associates each BGP neighbor relationship with individual VRFs. The routes that are received from each VRF routing protocol instance are inserted into the IP routing table that is contained within that VRF.

A per-VRF forwarding table, a FIB, is built from the per-VRF routing table and is used to forward all the packets that are received through the interfaces associated with the VRF.

https://t.me/learningnets



This figure illustrates the interactions between VRF instances of routing processes, VRF routing tables, and the global VPNv4 BGP routing process.

Example: BGP Route Propagation—Outbound

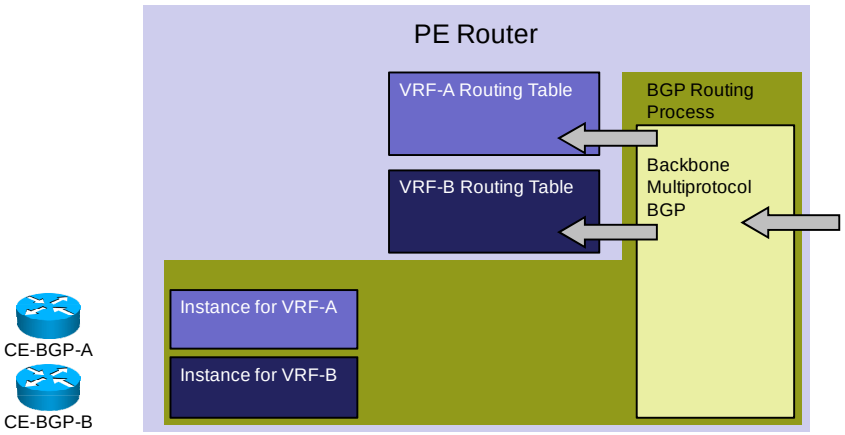
The BGP routes that are received from BGP-speaking CE routers are copied into the MP-BGP table for further propagation to other PE routers. This is the export process.

The IP prefixes are prepended with the RD, and the set of RTs (extended BGP communities) that are configured as export RTs for the VRF is attached to the resulting VPNv4 route.

Note	There are not separate per-VRF BGP and global MP-BGP tables in Cisco IOS Software.

https://t.me/learningnets

BGP Route Propagation—Inbound

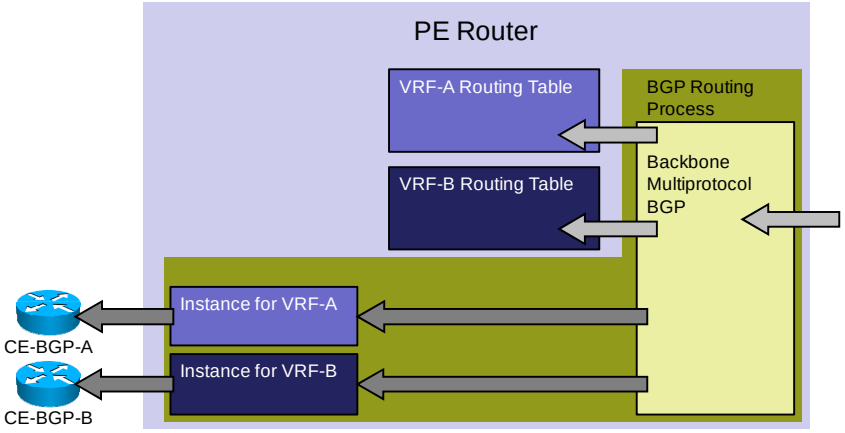


- VPNv4 prefixes are received from other PE routers.
- The VPNv4 prefixes are inserted into the proper VRF routing tables based on their route targets and the import route targets configured in VRFs.
- The route distinguisher is removed during this process.

© 2012 Cisco and/or its affiliates. All rights reserved.

As other PE routers start originating VPNv4 routes, the MP-BGP process in the PE router receives the routes. The routes are filtered, based on the RT attributes attached to them, and are inserted into the proper per-VRF IP routing tables, based on the *import RTs* that are configured for individual VRFs. The RD that was prepended by the originating PE router is removed before the route is inserted into the per-VRF IP routing table.

BGP Route Propagation—Inbound (Cont.)

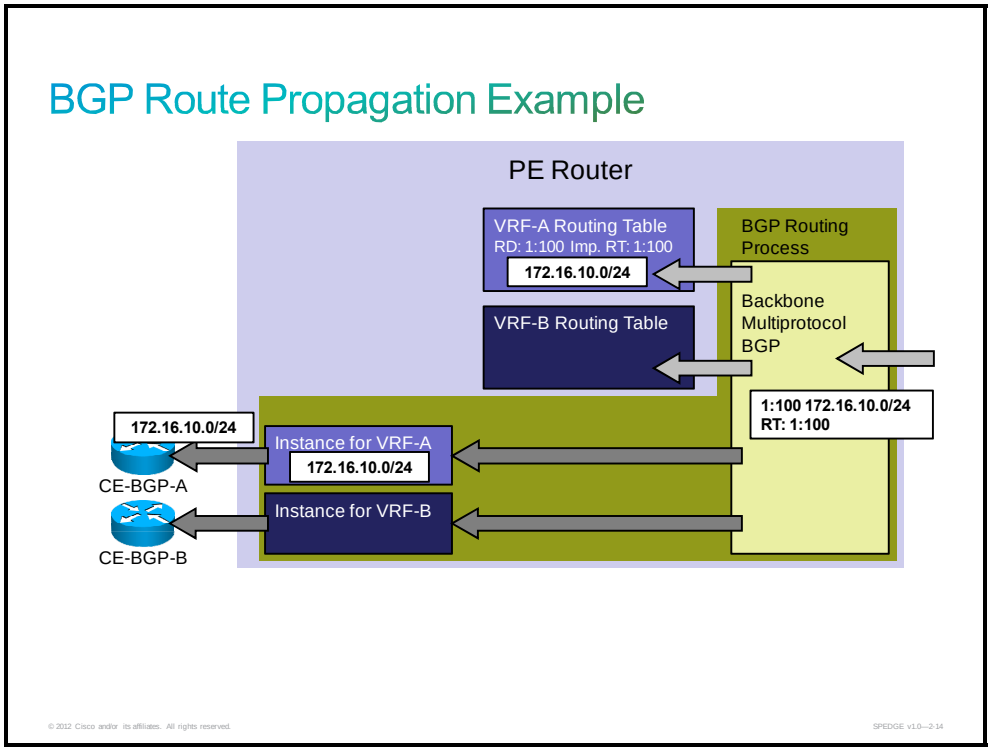


- Routes are received from backbone MP-BGP and imported into a VRF.
- IPv4 routes are forwarded to EBGP CE neighbors attached to that VRF.

© 2012 Cisco and/or its affiliates. All rights reserved.

The Multiprotocol Internal Border Gateway Protocol (MP-IBGP) VPNv4 routes that are received from other PE routers and selected by the import RTs of a VRF are automatically propagated as 32-bit IPv4 routes to all BGP-speaking CE neighbors of the PE router.

https://t.me/learningnets

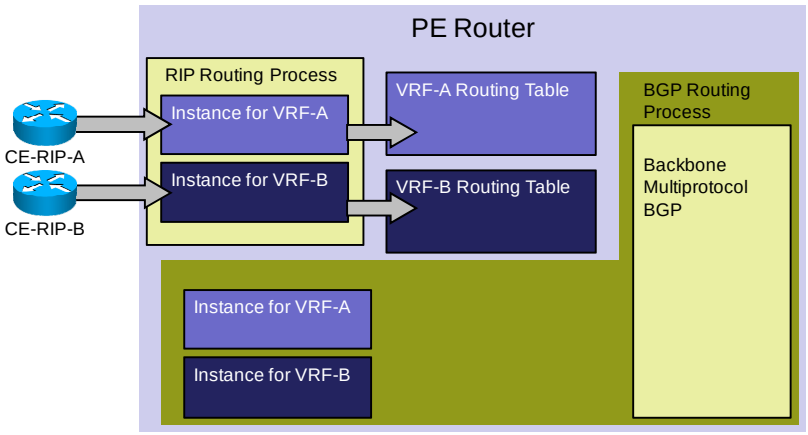


In this example, you see an incoming routing update for network 172.16.10.0/24 with RD 1:100. The routing update has also defined RT 1:100.

Because VRF-A has defined import RT 1:100, a routing update is inserted into the VRF-A routing table. A routing update is also sent to the CE-BGP-A router.

https://t.me/learningnets

Non-BGP Route Propagation—Outbound



- RIP-speaking CE routers announce their prefixes to the PE router via RIP.
- The instance of the RIP process associated with the VRF of the PE-CE interface collects the routes and inserts them into the VRF routing table.

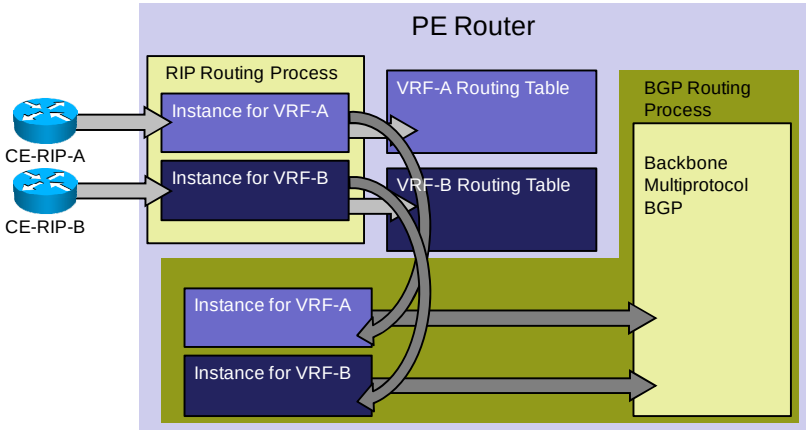
© 2012 Cisco and/or its affiliates. All rights reserved.

This example describes the outbound non-BGP route propagation process in an MPLS VPN implementation. The example describes RIP-speaking CE routers, but the process would be similar for other non-BGP protocols.

RIP-speaking CE routers identify the correct instance of RIP on the PE router when an inbound PE interface is associated with a VRF. This association allows CE routers to announce their networks to the appropriate per-VRF routing table.

https://t.me/learningnets

Non-BGP Route Propagation—Outbound (Cont.)



- The RIP routes entered in the VRF routing table are redistributed into BGP for further propagation into the MPLS VPN backbone.
- Redistribution between RIP and BGP has to be configured for proper MPLS VPN operation.

© 2012 Cisco and/or its affiliates. All rights reserved.

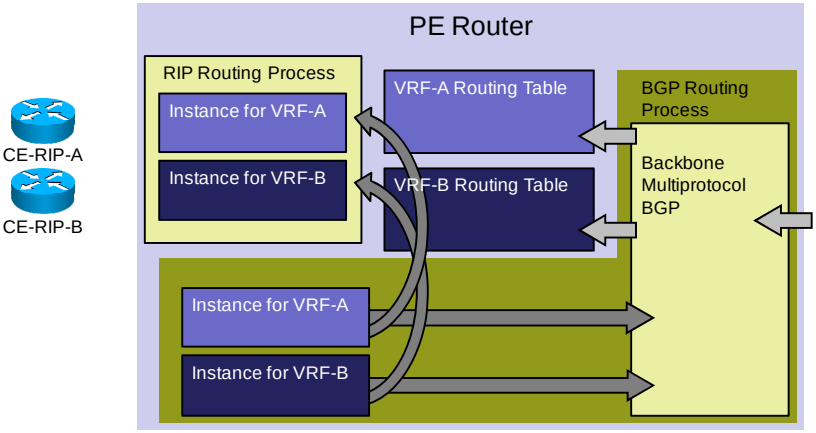
MP-BGP is used in the MPLS VPN backbone to carry VPN routes (prefixed with the RD) as 96-bit VPNv4 routes between the PE routers. The backbone BGP process looks exactly like a standard Internal Border Gateway Protocol (IBGP) setup from the perspective of the VRF. The per-VRF RIP routes therefore *must be redistributed* into the per-VRF instance of the BGP process to allow them to be propagated through the backbone MP-BGP process to other PE routers.

Caution Failure to redistribute non-BGP routes into the per-VRF instance of BGP is one of the most common MPLS VPN configuration errors.

If there is an overlap between an inbound RIP update and an inbound EBGP update, the standard route-selection mechanism (administrative distance) is used in the per-VRF IP routing table, and the EBGP route takes precedence over the RIP route. EBGP precedence results from the fact that the administrative distance of EBGP routes (20) is better than the administrative distance of RIP routes (120).

https://t.me/learningnets

Non-BGP Route Propagation—Inbound

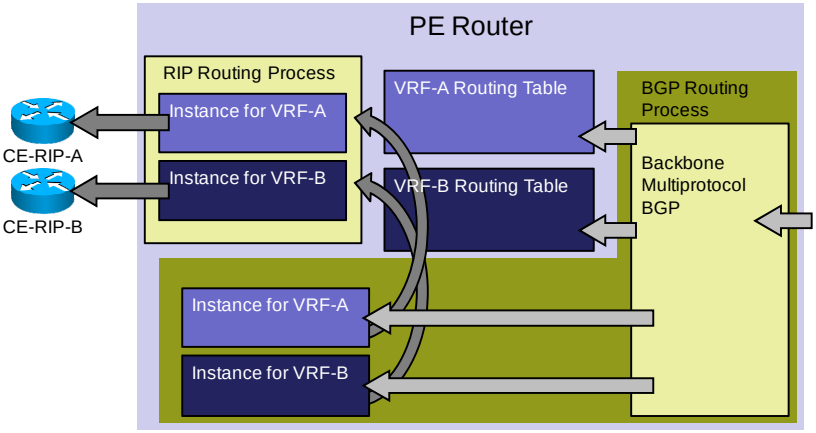


- The RIP routes entered in the VRF routing table are redistributed into BGP for further propagation into the MPLS VPN backbone.
- Redistribution between RIP and BGP has to be configured for proper MPLS VPN operation.

© 2012 Cisco and/or its affiliates. All rights reserved.

The MP-IBGP routes, although they are inserted in the per-VRF IP routing table, are *not* propagated to RIP-speaking CE routers automatically. To propagate these MP-IBGP routes to the RIP-speaking CE routers, you must manually configure redistribution between the per-VRF instance of BGP and the per-VRF instance of RIP.

Non-BGP Route Propagation—Inbound (Cont.)

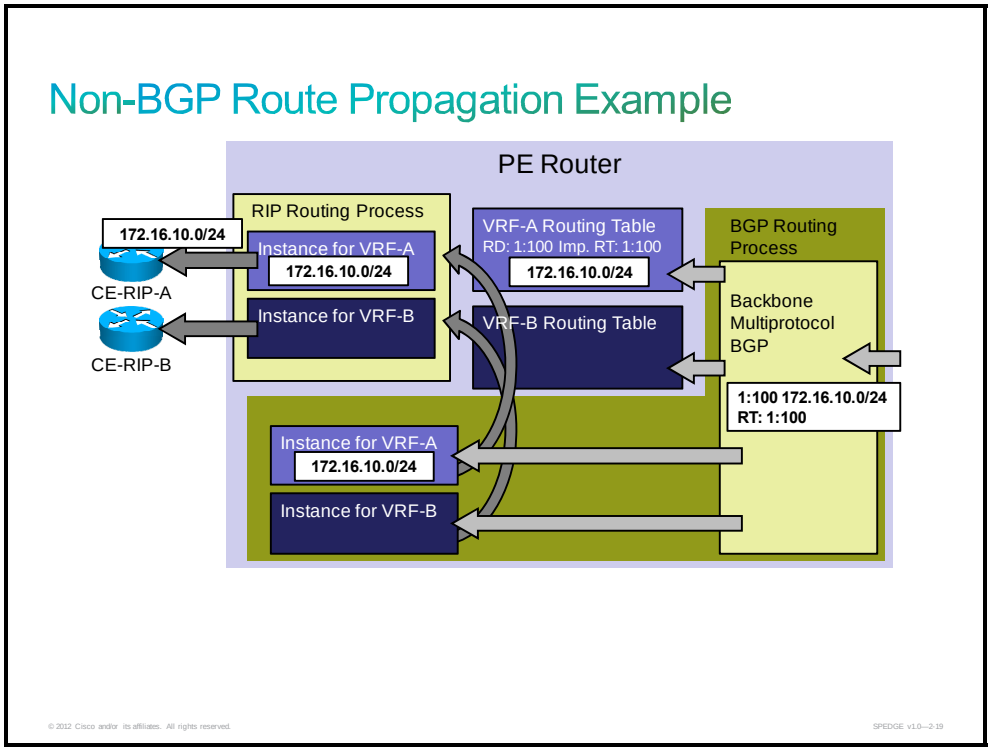


- Routes redistributed from BGP into a VRF instance of RIP are sent to RIP-speaking CE routers.

© 2012 Cisco and/or its affiliates. All rights reserved.

When the IBGP routes from the per-VRF IP routing table are successfully redistributed into the per-VRF instance of the RIP process, the RIP process announces these routes to RIP-speaking CE routers, thus achieving transparent end-to-end connectivity between the CE routers.

https://t.me/learningnets

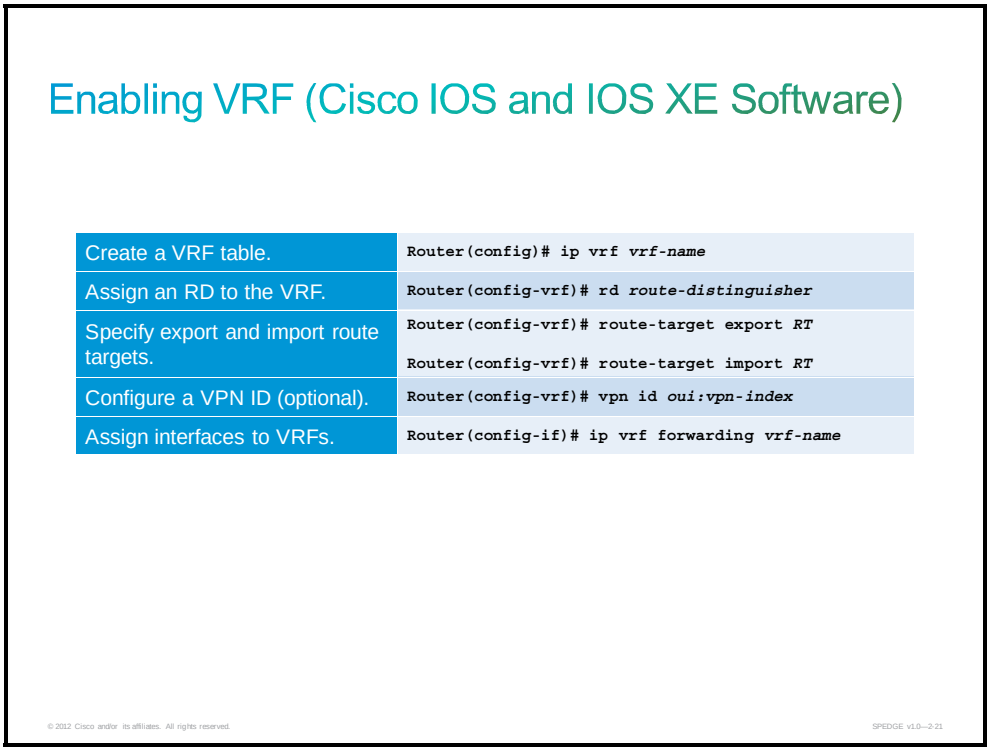


In this example, you can see a routing update for network 172.16.10.0/24. A route is inserted in the VRF-A routing table. The routing update is processed by the VRF-A instance of the BGP routing process and redistributed to the VRF-A instance of the RIP routing process. The routing update is then sent to the CE-RIP-A router by the RIP routing process.

https://t.me/learningnets

Enabling VRF

This topic describes how to enable VRF.



Configuring a VRF table and starting deployment of an MPLS VPN service for a customer on Cisco IOS and IOS XE platform consists of these four mandatory steps:

1. Create a new VRF table.
2. Assign a unique RD to the VRF.

Note	You must assign a unique RD to every VRF created in a PE router. The same RD *might* be used in multiple PE routers, based on customer connectivity requirements. The same RD *should* be used on all PE routers for simple VPN service.

3. Specify import and export RTs for the VRF.

Note	Import and export RTs should be equal to the RD for simple VPN service.

4. Assign interfaces to the VRF.

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Enabling VRF (Cisco IOS XR Software)

Create a VRF table.	`RP/0/RP0/CPU0:router(config)# vrf vrf-name`
Enter VRF address family configuration mode for the IPv4 address family.	`RP/0/RP0/CPU0:router(config-vrf)# address-family ipv4 unicast`
Specify import route targets.	`RP/0/RP0/CPU0:router(config-vrf-af)# import route-target [as-number:nn | ip-address:nn]`
Specify export route targets.	`RP/0/RP0/CPU0:router(config-vrf-af)# export route-target [as-number:nn | ip-address:nn]`
Assign interfaces to VRFs.	`RP/0/RP0/CPU0:router(config-if)# vrf vrf-name`

© 2012 Cisco and/or its affiliates. All rights reserved.

Configuring a VRF table on Cisco IOS XR devices is somewhat different from using Cisco IOS and IOS XE Software. Basic configuration consists of these four mandatory steps:

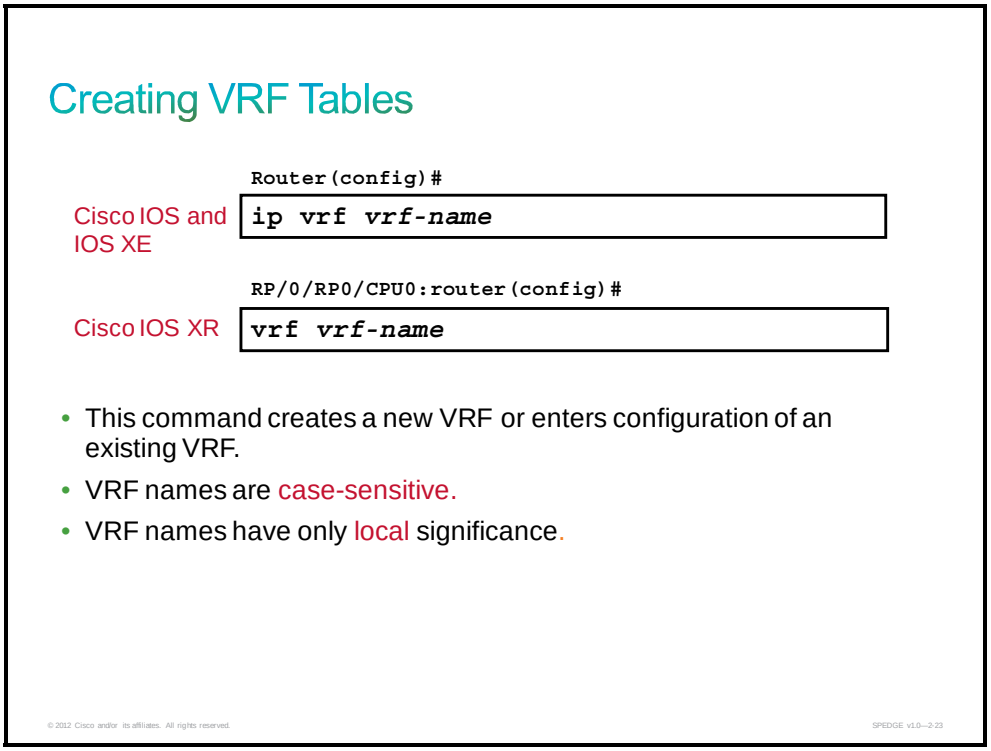
Step 1 Create a new VRF table.

Step 2 Enter the IPv4 unicast address family configuration.

Step 3 Specify import and export RTs for the VRF.

Step 4 Assign interfaces to the VRF.

https://t.me/learningnets



To configure a VRF routing table on the Cisco IOS and IOS XE platforms, use the **ip vrf** command in global configuration mode. To remove a VRF routing table, use the **no** form of this command.

- **ip vrf** *vrf-name*
- **no ip vrf** *vrf-name*

To configure a VRF routing table on the Cisco IOS XR platform, use the **vrf** command in global configuration mode. To remove a VRF routing table, use the **no** form of this command.

- **vrf** *vrf-name*
- **no vrf** *vrf-name*

No VRFs are defined by default. No import or export lists are associated with a VRF. No route maps are associated with a VRF.

https://t.me/learningnets

Assigning RDs

- This command assigns a route distinguisher to a VRF.
- A VRF is not operational unless you configure an RD.
- You can use the ASN:nn or A.B.C.D:nn format for RD.
- Each VRF in a PE router must have a unique RD.

Cisco IOS and IOS XE configuration

RD is configured under VRF configuration area

```
Router(config)#ip vrf vrf-name
Router(config-vrf)#rd route-distinguisher
```

Cisco IOS XR configuration

RD is configured under BGP configuration area

```
Router(config)#router bgp AS
Router(config-bgp)#vrf vrf-name
Router(config-bgp-vrf)#rd route-distinguisher
```

© 2012 Cisco and/or its affiliates. All rights reserved.

To create routing and forwarding tables for a VRF on the Cisco IOS and IOS XE platforms, use the **rd** command in VRF configuration submode: **rd** *route-distinguisher*. The *route-distinguisher* parameter adds an 8-byte value to an IPv4 prefix to create a VPNv4 prefix.

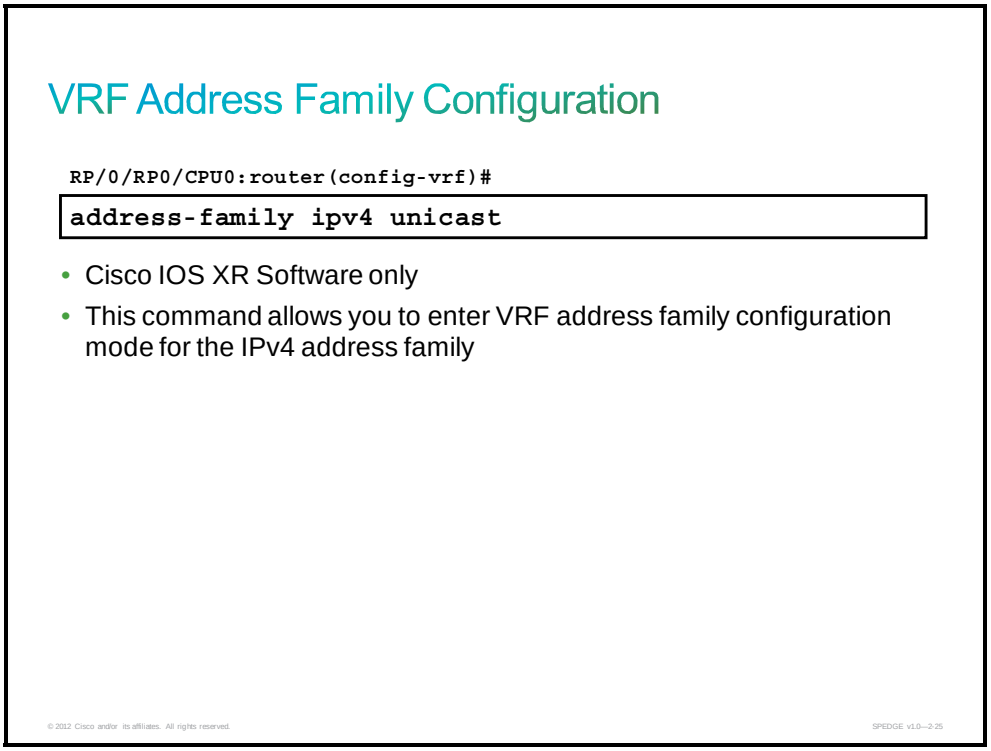
The RD can be specified in one of these two formats:

- 16-bit autonomous system (AS) number followed by a 32-bit decimal number (ASN:nn)
- 32-bit IP address followed by a 16-bit decimal number (A.B.C.D:nn)

There is no default for this command. An RD must be configured for a VRF table to be functional.

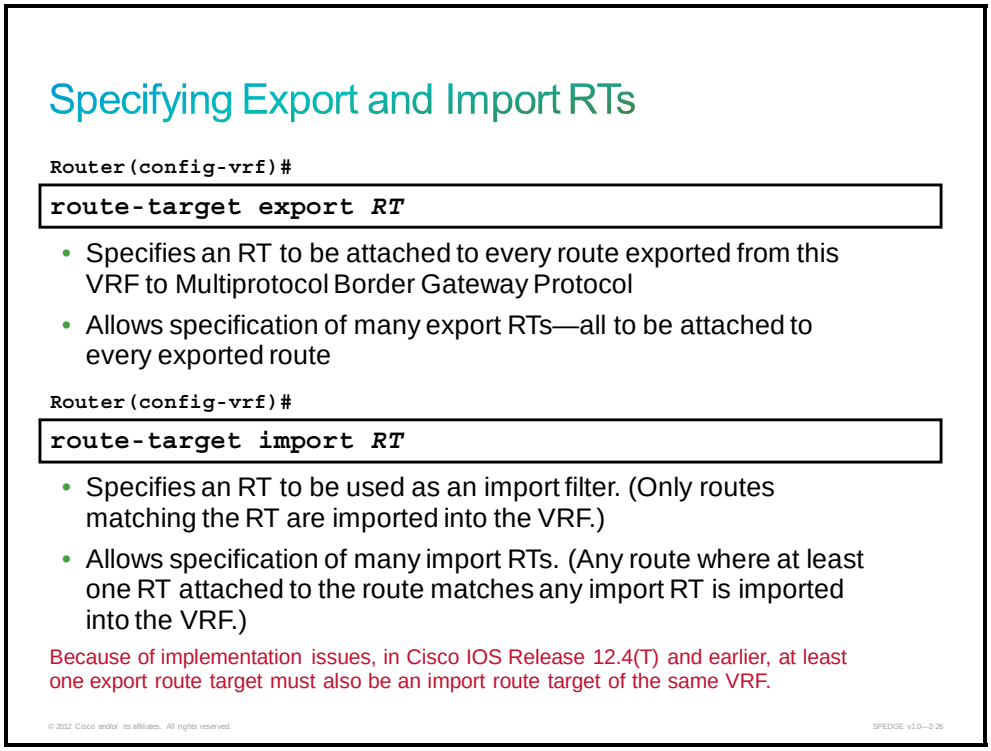
To create routing and forwarding tables for a VRF on a Cisco IOS XR operating system, use the **rd** command in the BGP configuration area in VRF configuration submode.

https://t.me/learningnets



To create routing and forwarding tables for a VRF on the Cisco IOS XR platform, you must first enter VRF address family configuration submode using the **address-family ipv4 unicast** command. Address families are used within VRF configuration mode to control import and export policies.

https://t.me/learningnets



To create an RT extended community for a VRF on the Cisco IOS and IOS XE platforms, use the **route-target** command in VRF submode. To disable the configuration of an RT community option, use the **no** form of this command.

- **route-target** {**import** | **export** | **both**} *route-target-ext-community*
- **no route-target** {**import** | **export** | **both**} *route-target-ext-community*

This table describes the parameters for the **route-target** command.

Syntax Description

Parameter	Description
import	Imports routing information from the target VPN extended community
export	Exports routing information to the target VPN extended community
both	Sets the value to be used by both the import and export process to the value that is indicated in the *route-target-ext-community* field
route-target-ext-community	Adds the route target extended community attributes to the VRF list of import, export, or both (import and export) route target extended communities

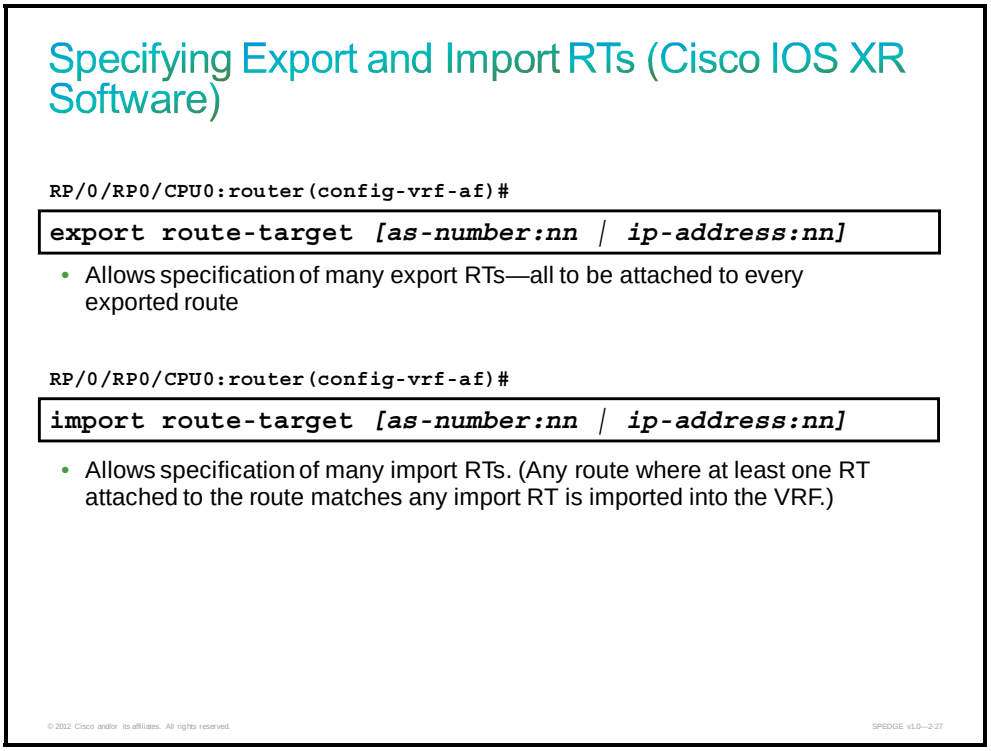
Similar to RDs, RTs can be specified in one of these two formats:

- 16-bit AS number followed by a 32-bit decimal number (ASN:nn)
- 32-bit IP address followed by a 16-bit decimal number (A.B.C.D:nn)

A VRF has no RT extended community attributes associated with it until they are specified by the **route-target** command.

Whenever an RT is both an import RT and an export RT for a VRF, you can use the **route-target both** command to simplify the configuration.

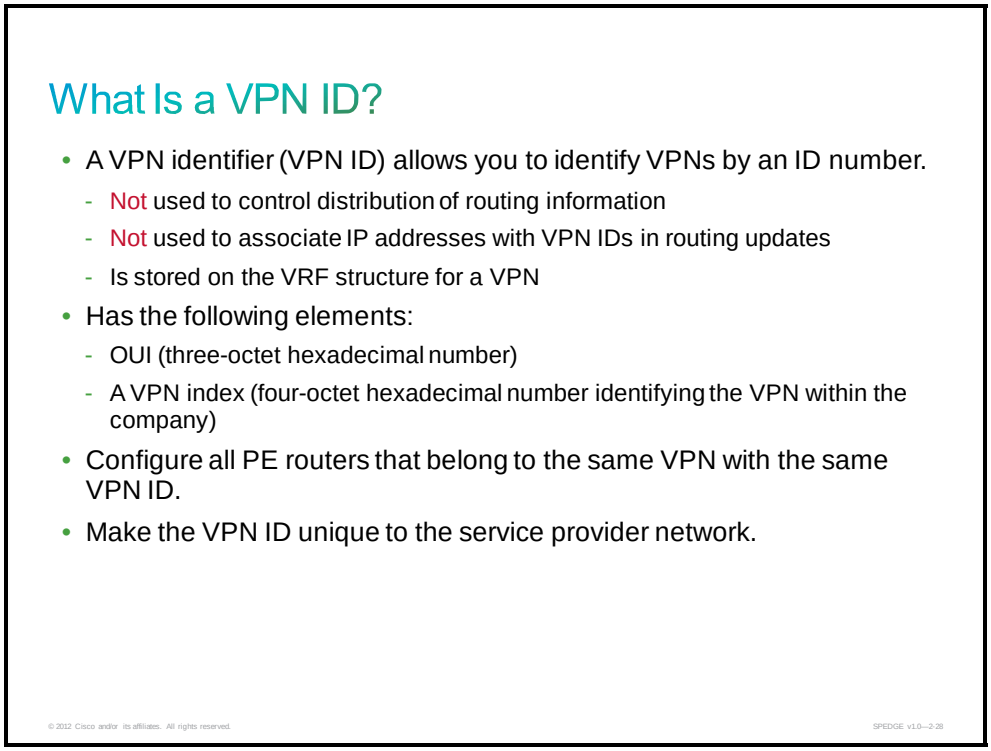
https://t.me/learningnets



The Cisco IOS XR **export route-target** command associates the local VPN with an RT. When the route is advertised to other PE routers, the export RT is sent along with the route as an extended community.

The **import route-target** command allows exported VPN routes to be imported into the VPN if one of the RTs of the exported route matches one of the local VPN import RTs.

https://t.me/learningnets



The MPLS VPN ID is an optional feature that allows you to identify VPNs by a VPN identification number. The MPLS VPN ID feature is not used to control the distribution of routing information or to associate IP addresses with MPLS VPN ID numbers in routing updates.

You can configure all the PE routers that belong to the same VPN with the same VPN ID. Make sure that the VPN ID is unique to the service provider network.

The VPN ID is stored in the corresponding VRF structure for the VPN. To ensure that the VPN has a consistent VPN ID, assign the same VPN ID to all the routers in the service provider network that service that VPN.

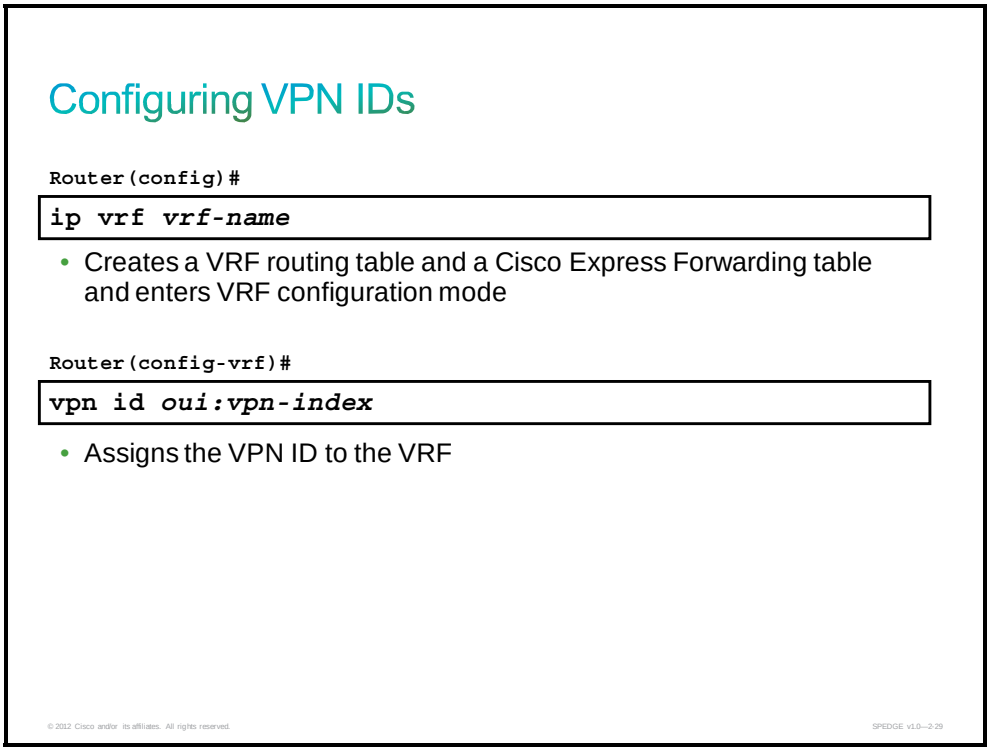
Each VPN ID that is defined by RFC 2685 consists of these elements:

- An Organizationally Unique Identifier (OUI), a three-octet hexadecimal number that is assigned by the IEEE
- A VPN index, a four-octet hexadecimal number that identifies the VPN within the company

A VPN ID is useful for remote access applications, such as RADIUS and DHCP, which can use the MPLS VPN ID to identify a VPN. RADIUS can use the VPN ID to assign dial-in users to the proper VPN, based on the authentication information of each user.

Note	You can use a VRF name (a unique ASCII string) to reference a specific VPN that is configured in the router. Alternatively, you can use a VPN ID to identify a particular VPN that is configured in the router. The VPN name is not affected by the VPN ID configuration.

https://t.me/learningnets



To assign a VPN ID to a VRF, use the **vpn id** command in VRF configuration submode. To disable the configuration of an RT community option, use the **no** form of this command.

- **vpn id** *oui:vpn-index*
- **no vpn id** *oui:vpn-index*

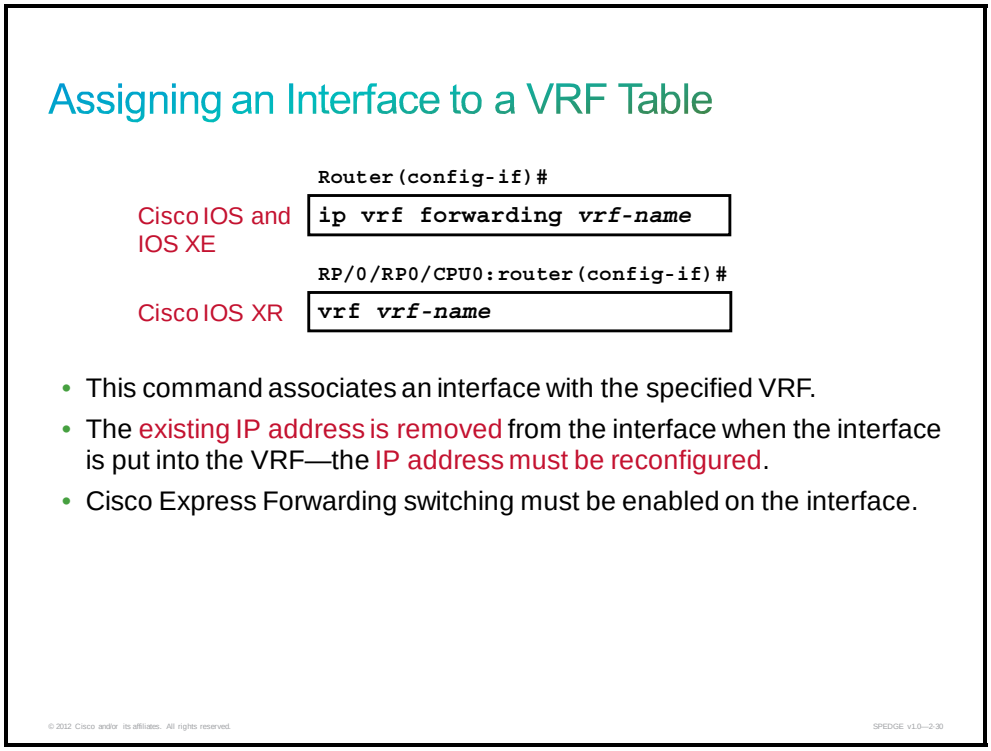
This table describes the parameters for the **vpn id** command.

Syntax Description

Parameter	Description
oui	Identifies the OUI, which is restricted to three octets and is followed by a colon
vpn-index	Identifies the VPN within the company and is restricted to four octets

Each VRF configured in a PE router can have a VPN ID configured. Configure all the PE routers that belong to the same VPN with the same VPN ID. The VPN ID should be unique to the service provider network.

https://t.me/learningnets



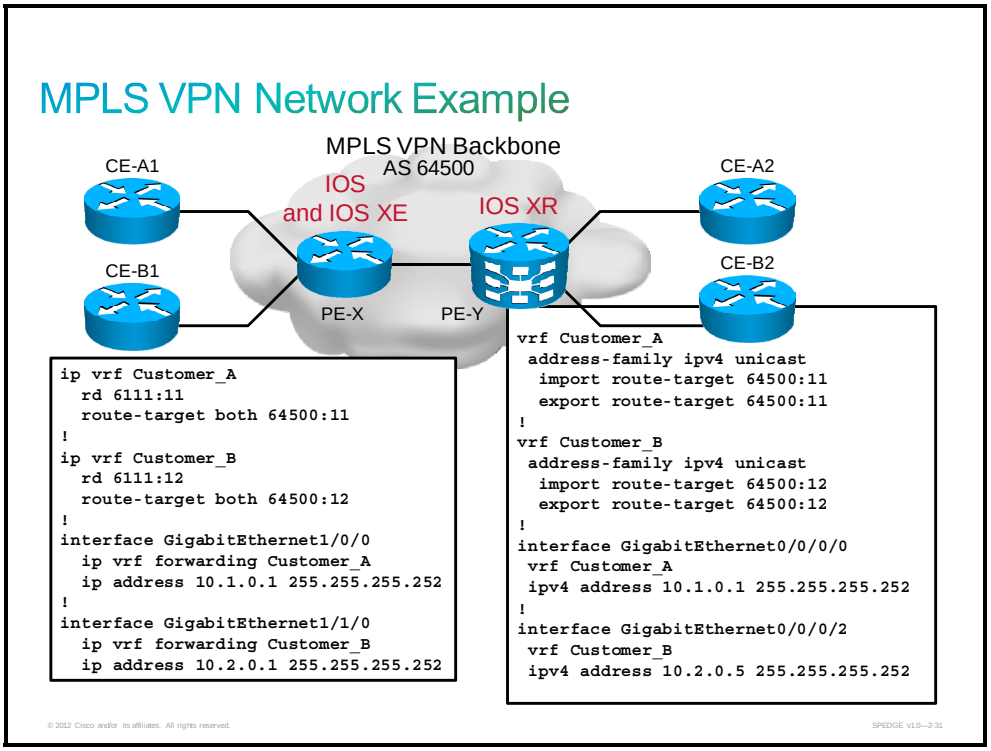
To associate a VRF with an interface or subinterface, use the **ip vrf forwarding** command on the Cisco IOS and IOS XE platforms or the **vrf** command on the Cisco IOS XR platform in interface configuration mode. To disassociate a VRF, use the **no** form of this command.

Note	You must remove IPv4 and IPv6 addresses from an interface before assigning, removing, or changing its VRF. If you do not, any attempt to change the VRF on an IP interface is rejected.

After local interfaces are bound to the VRF, the interfaces appear in the routing display of the VRF table.

Note	When an interface is configured with a particular VRF, its IP address is removed from the interface and from the global routing table. This action is based on the assumption that the address is not valid across multiple routing tables and that the address should be reconfigured after the interface is associated to a VRF.

https://t.me/learningnets



To illustrate the use of MPLS VPN configuration commands, you can look at a configuration of the PE routers in a sample network.

The figure shows configuration of the PE routers in a sample network with two VPN customers. The PE-X router is running Cisco IOS or IOS XE Software. The PE-Y router is running Cisco IOS XR Software.

The configuration steps that you perform on the PE router are as follows:

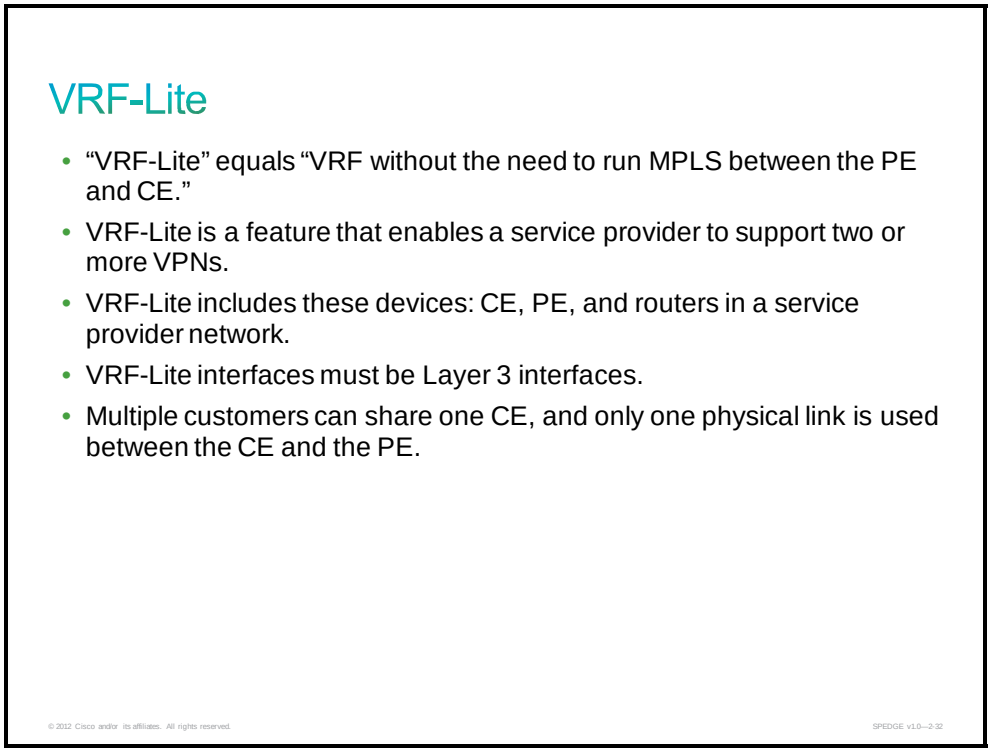
Step 1 Configure VRFs for customer A and customer B.

Step 2 Assign RDs and RTs to the VRFs.

Only one RD per customer is used on all PE routers in this MPLS VPN backbone, because these customers require only simple VPN connectivity. To simplify the configuration and troubleshooting process, the RTs are made equal to the RDs.

Step 3 Assign PE-CE interfaces to individual VRFs.

https://t.me/learningnets



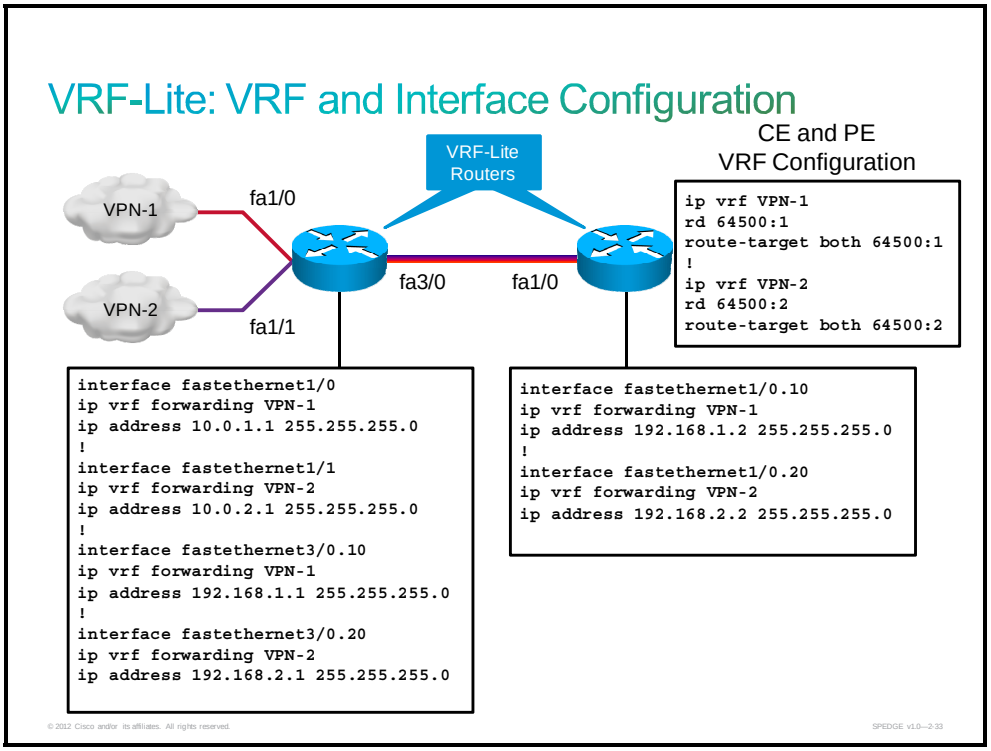
Multi-VRF Customer Edge (VRF-Lite) provides the ability to configure and maintain more than one VRF instance within the same CE router.

VRF-Lite uses input interfaces to distinguish routes for different VPNs and forms VRF tables by associating one or more Layer 3 interfaces with each VRF. Interfaces in a VRF can be either physical, such as Ethernet ports, or logical, such as VLAN switch virtual interfaces (SVIs). However, a Layer 3 interface cannot belong to more than one VRF at any time. The VRF-Lite feature thus allows an operator to support two or more routing domains on a CE router, with each routing domain having its own set of interfaces and its own set of routing and forwarding tables.

VRF-Lite includes these devices:

- **Customer edge devices:** CE devices provide customer access to the service provider network over a data link to one or more PE routers. The CE device advertises the local routes of the site to the PE router and learns the remote VPN routes from it. A Cisco Catalyst 4500 Series Switch can be a CE.
- **Provider edge routers:** PE routers exchange routing information with CE devices by using static routing or a routing protocol such as BGP, RIPv1, or RIPv2. The PE router is only required to maintain VPN routes for the VPNs to which it is directly attached, eliminating the need for the PE router to maintain all of the service provider VPN routes. Each PE router maintains a VRF for each of its directly connected sites. Multiple interfaces on a PE router can be associated with a single VRF if all of these sites participate in the same VPN. Each VPN is mapped to a specified VRF. After learning local VPN routes from CE devices, a PE router exchanges VPN routing information with other PE routers by using IBGP.

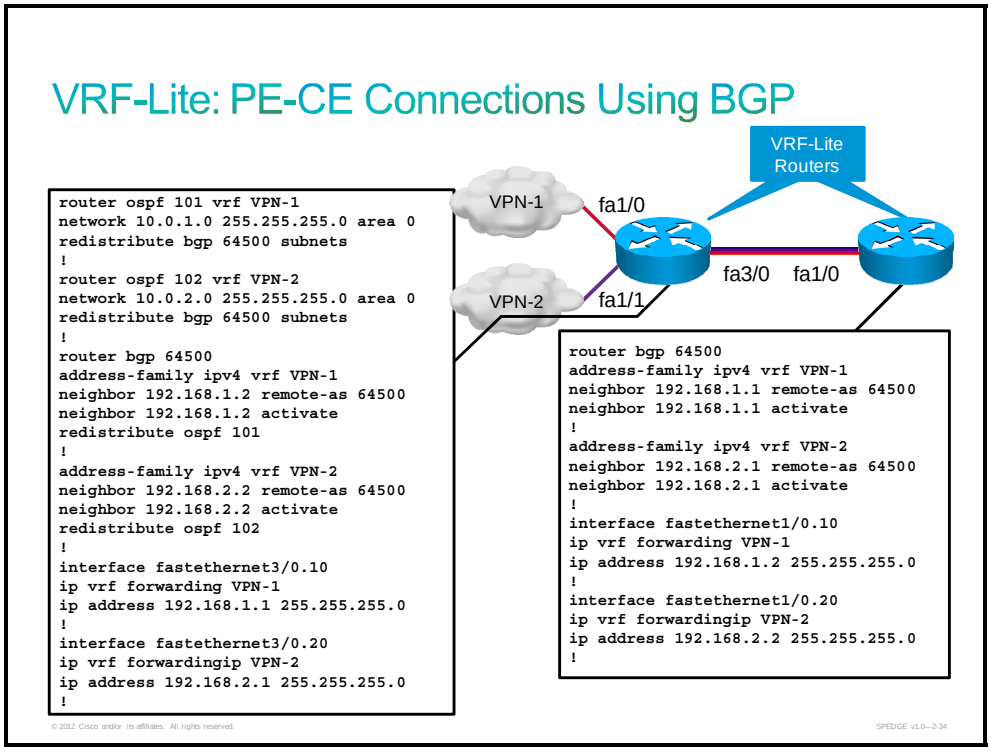
https://t.me/learningnets



With VRF-Lite, multiple customers can share one CE, and only one physical link is used between the CE and the PE. The shared CE maintains separate VRF tables for each customer and switches or routes packets for each customer, based on its own routing table. VRF-Lite extends limited PE functionality to a CE device, giving it the ability to maintain separate VRF tables to extend the privacy and security of a VPN to the branch office.

To illustrate VRF-Lite configuration, you can look at a configuration of the CE and PE routers in a sample network. First, VRFs must be configured on both the PE and the CE routers. Additionally, you must specify the Layer 3 interface to be associated with the VRF and associate the VRF with the Layer 3 interface.

https://t.me/learningnets



Most routing protocols can be used between the CE and the PE: BGP, OSPF, EIGRP, RIP, and static routing. However, EBGP is recommended:

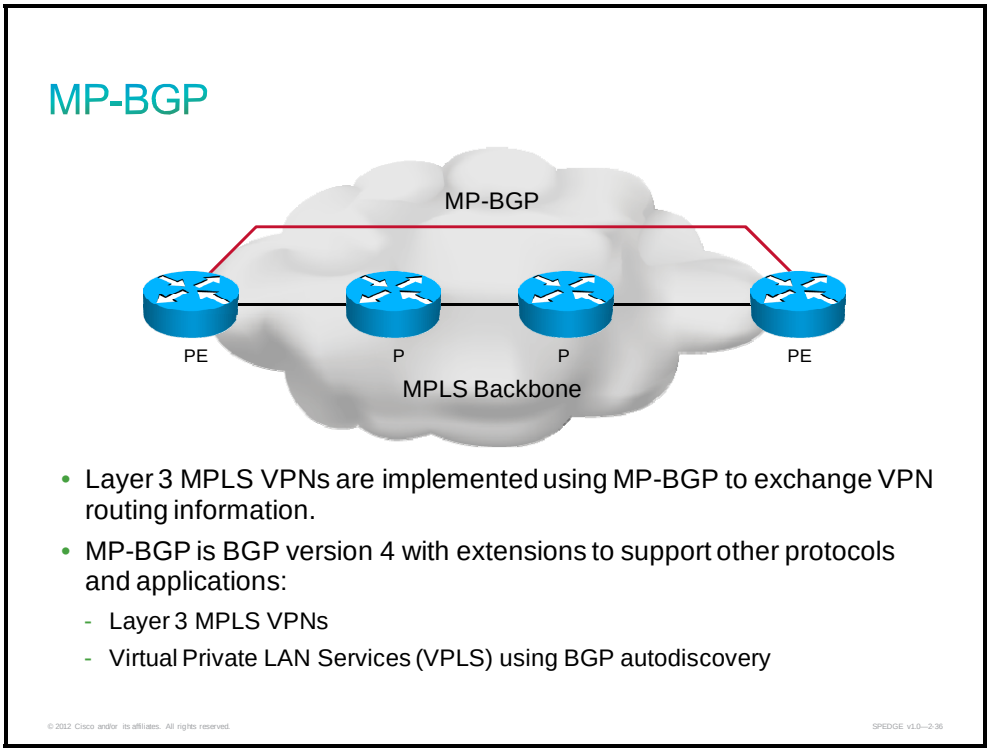
- BGP does not require more than one algorithm to communicate with a multitude of CEs.
- BGP is designed to pass routing information between systems that are run by different administrations.
- BGP makes it easy to pass attributes of the routes to the CE.

Furthermore, when BGP is used as the routing protocol, it can also be used to manage the MPLS label exchange between the PE and CE devices. By contrast, if OSPF, EIGRP, RIP, or static routing is used, Label Distribution Protocol (LDP) must be used to signal labels.

https://t.me/learningnets

Enabling MP-BGP

MP-BGP is responsible for allocating labels for VPN routes and advertising them to other edge routers. This topic describes how to enable MP-BGP.



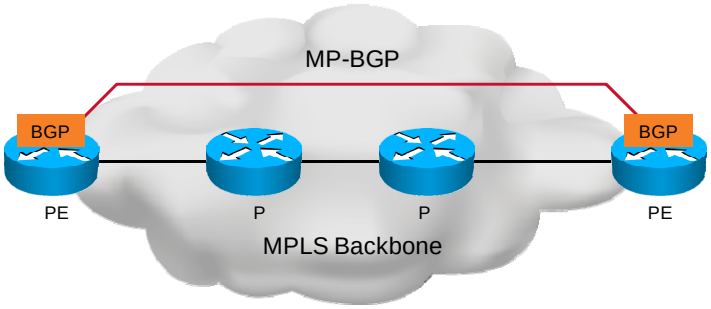
VPNs based on MPLS require an additional VPN label to distinguish between potentially overlapping prefixes belonging to different VPNs. MP-BGP is BGP version 4 with additional attributes to support the exchange of VPN prefixes.

Virtual Private LAN Services (VPLS) can also be implemented using the BGP autodiscovery feature to simplify the management of VPLS.

The figure shows an end-to-end MP-IBGP session. This figure is a simplified representation of the BGP capability to propagate VPN routing information between edge label switch routers (LSRs). In real environments with many more PE routers, a route reflector would be used between the edge routers, although that addition would not alter the operation significantly.

https://t.me/learningnets

MP-BGP Configuration Tasks



- MP-BGP must be configured on edge routers only.
- Support for MPLS VPNs must be enabled.
- Steps required:
  - Add address family vpnv4
  - Activate neighbor in address family vpnv4
- Optional configuration settings

© 2012 Cisco and/or its affiliates. All rights reserved.

MP-BGP extension is used in the MPLS world to relay VPN information between two edge routers. The RD is a 64-bit value that is used to mark prefixes and to separate different customers.

VPN support in BGP is enabled by configuring a VPNv4 address family. This allows MP-BGP neighbor sessions to be established independently from existing IPv4 BGP sessions. These VPNv4 adjacencies are used to relay VPN prefixes together with 64-bit extended communities, where the RT value is stored. The total length of the VPNv4 address is thus 96 bits.

Configuring the VPNv4 address family and activating neighbors in it is the minimum required configuration. Optionally, fine-tuning can be performed by adjusting the BGP timers.

https://t.me/learningnets

BGP Address Families

- The BGP process in an MPLS VPN-enabled router performs three separate tasks:
  - Global BGP routes (Internet routing) are exchanged as in a traditional BGP setup.
  - VPNv4 prefixes are exchanged through MP-BGP.
  - VPN routes are exchanged with CE routers through per-VRF External Border Gateway Protocol sessions or through route redistribution.
- Address families (routing protocol contexts) are used to configure these three tasks in the same BGP process.

© 2012 Cisco and/or its affiliates. All rights reserved.

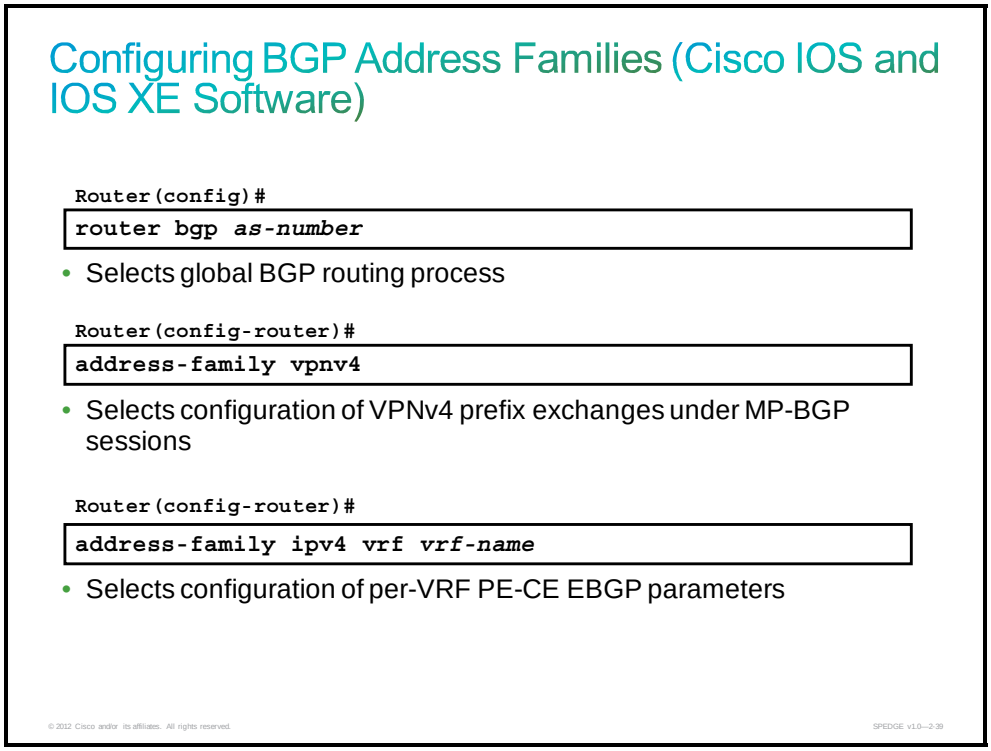
Independently from the MPLS VPN architecture, the PE router can use BGP IPv4 route updates to receive and propagate Internet routes in situations where the PE routers are also used to provide Internet connectivity to customers.

The MPLS VPN architecture uses the BGP routing protocol in these two ways:

- VPNv4 routes are propagated across an MPLS VPN backbone using MP-BGP between the PE routers.
- BGP can be used as the PE-CE routing protocol to exchange VPN routes between the PE routers and the CE routers.

All three route-exchange mechanisms take place in one BGP process (because only one BGP process can be configured per router). The routing protocol contexts (called address families from the router configuration perspective) are used to configure all three independent route-exchange mechanisms.

https://t.me/learningnets



To configure the BGP routing process, use the **router bgp** command in global configuration mode. To remove a routing process, use the **no** form of this command.

- **router bgp** *as-number*
- **no router bgp** *as-number*

Use the **address-family** command in router configuration mode to select the routing context that you would like to configure:

- Internet routing (global IP routing table) is the default address family that you configure when you start configuring the BGP routing process.
- To configure MP-BGP sessions between the PE routers, use the **address-family vpnv4** command.
- To configure BGP between the PE routers and the CE routers within individual VRF tables, use the **address-family ipv4** *vrf vrf-name* command.

To enter address-family submode for configuring routing protocols, such as BGP, RIP, and static routing, use the **address-family** command in global configuration mode. To disable address-family submode for configuring routing protocols, use the **no** form of this command.

- VPNv4 unicast: **address-family vpnv4** [**unicast**]
  Configures sessions that carry customer VPNv4 prefixes, each of which has been made globally unique by adding an 8-byte RD
- IPv4 unicast: **address-family ipv4** [**unicast**]
  Configures sessions that carry standard IPv4 address prefixes
- IPv4 unicast: **address-family ipv4** [**unicast**] **vrf** *vrf-name*
  Specifies the name of a VPN VRF to associate with submode commands

https://t.me/learningnets

Configuring BGP Address Families (Cisco IOS XR Software)

```
RP/0/RP0/CPU0:router(config) #
router bgp as-number
```

- Selects global BGP routing process

```
RP/0/RP0/CPU0:router(config-bgp) #
address-family vpnv4 unicast
```

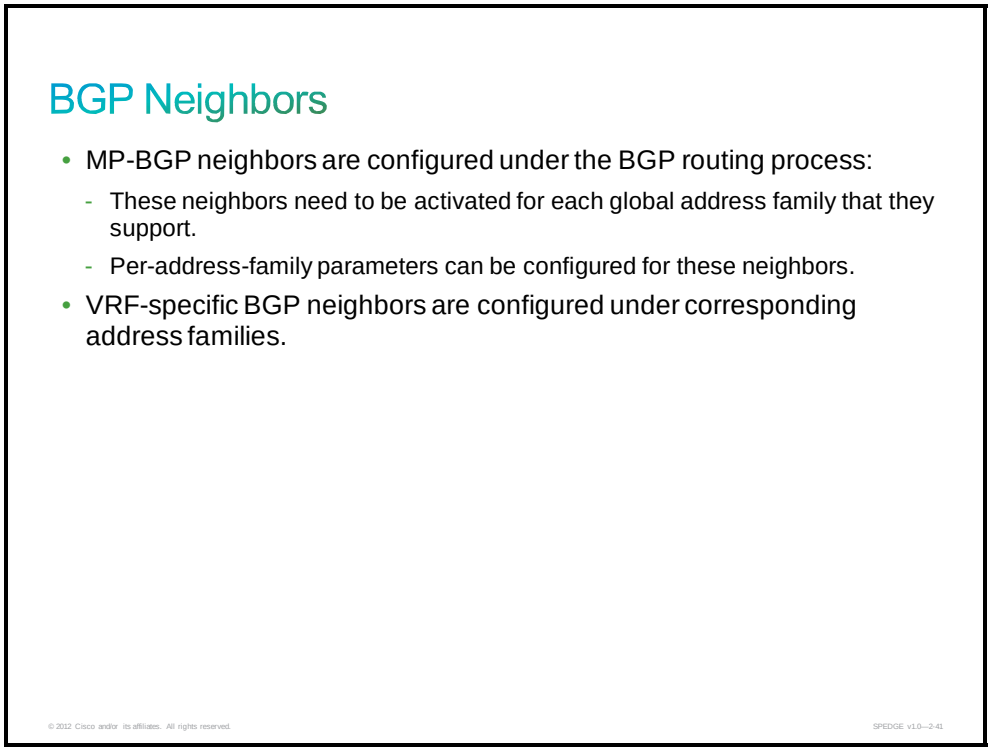
- Configures VPNv4 prefix

© 2012 Cisco and/or its affiliates. All rights reserved.

Similar to Cisco IOS and IOS XE Software, on Cisco IOS XR Software, use the **router bgp** command in global configuration mode.

The VPNv4 address family is configured in the BGP section using the **address-family vpnv4 unicast** command. Afterwards, it will be applied in the neighbor configuration block.

https://t.me/learningnets



MPLS VPN architecture defines these two types of BGP neighbors:

- Global BGP neighbors (other PE routers) with which the PE router can exchange multiple types of routes. (These neighbors are defined in the global BGP definition and need to be activated only for individual address families.)
- Per-VRF BGP neighbors. (These neighbors are the CE routers.)

https://t.me/learningnets

Configuring MP-IBGP on Cisco IOS and IOS XE Software

```
Router(config)#
router bgp as-number
 neighbor ip-address remote-as as-number
 neighbor ip-address update-source interface-type
      interface-number
```

- All MP-BGP neighbors have to be configured under the global BGP routing configuration.
- MP-IBGP sessions have to run between loopback interfaces.

```
Router(config-router)#
address-family vpnv4
```

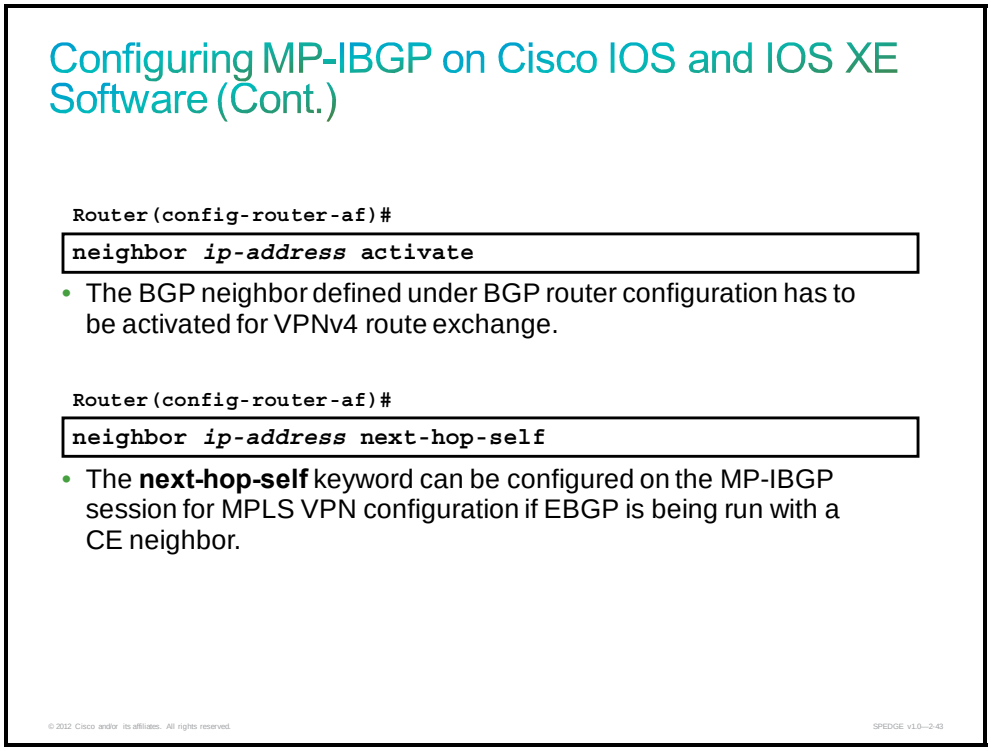
- This command starts configuration of MP-BGP routing for VPNv4 route exchange.
- The parameters that apply only to MP-BGP exchange of VPNv4 routes between already configured IBGP neighbors are configured under this address family.

© 2012 Cisco and/or its affiliates. All rights reserved.

The initial commands that are needed to configure an MP-IBGP session between PE routers are as follows:

- The **neighbor** *ip-address* **remote-as** *as-number* command configures the neighboring PE router.
- The **neighbor** *ip-address* **update-source** *interface-type interface-number* command configures the source address that is used for the TCP session carrying BGP updates and the IP address that is used as the BGP next hop for VPNv4 routes.
- The **address-family vpnv4** command allows you to enter VPNv4 configuration mode, where additional VPNv4-specific parameters must be configured on the BGP neighbor.

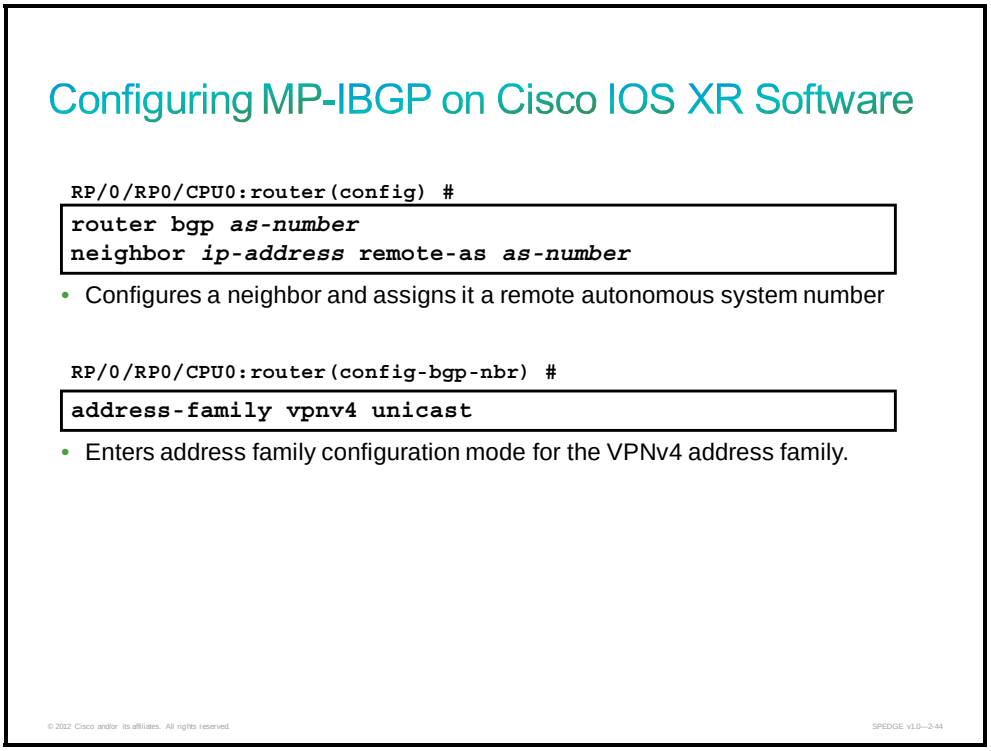
https://t.me/learningnets



After you define the remote PE router as a global BGP neighbor, you must activate it for VPNv4 route exchange. To enable the exchange of information with a BGP neighboring router, use the **neighbor activate** command in router configuration mode. The exchange of addresses with neighbors is enabled by default for the IPv4 address family. For all other address families, address exchange is disabled by default. You can explicitly activate the default command by using the appropriate address family submode.

To enable next-hop processing of BGP updates on the router, use the **neighbor next-hop-self** command in router configuration mode. This command is useful in unmeshed networks (such as Frame Relay or X.25) where BGP neighbors might not have direct access to all other neighbors on the same IP subnet. If you specify a BGP peer group by using the *peer-group-name* argument, all the members of the peer group inherit the characteristic that is configured with this command. Specifying the command with an IP address overrides the value that is inherited from the peer group.

https://t.me/learningnets

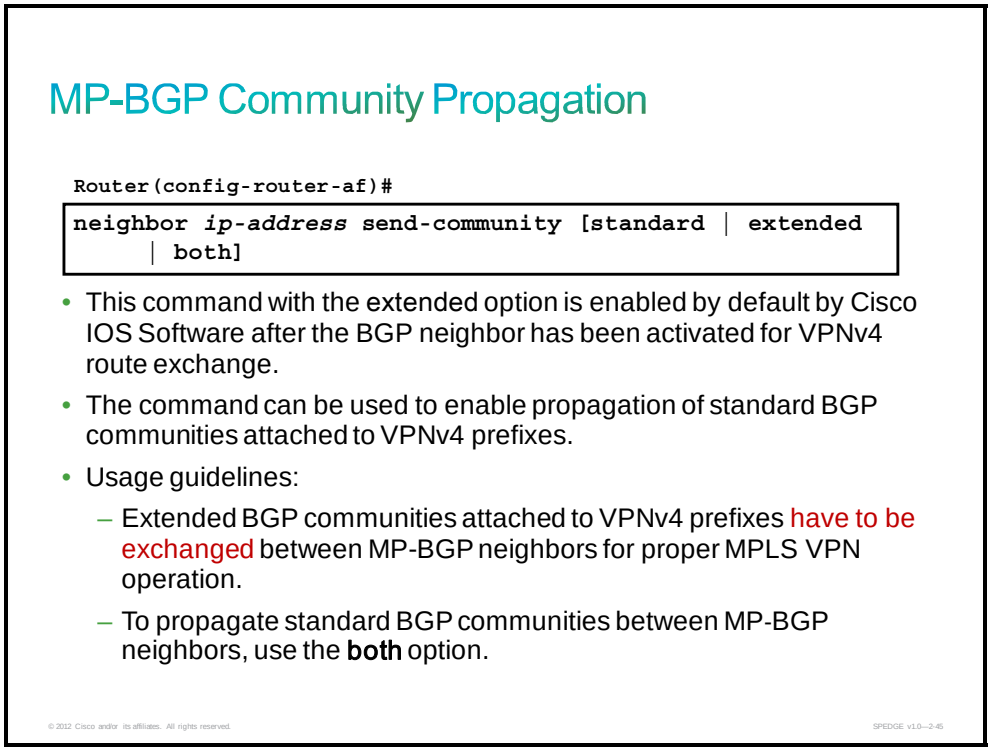


To configure an MP-IBGP neighbor on devices running Cisco IOS XR Software, enter BGP configuration mode using the command **router bgp** *as-number*.

To add a new BGP neighbor, use the command **neighbor** *ip-address* **remote-as** *as-number*.

In each neighbor configuration area, enable the neighbor for the specific address family.

https://t.me/learningnets



MPLS VPN architecture introduced the extended community BGP attribute. BGP still supports the standard community attribute, which has not been superseded by extended communities. The default community propagation behavior for standard BGP communities has not changed. Community propagation still must be configured manually. Extended BGP communities are propagated by default because their propagation is mandatory for successful MPLS VPN operation.

The **neighbor send-community** command was extended to support standard and extended communities. Use this command to configure propagation of standard and extended communities if your BGP design relies on use of standard communities. An example would be to propagate quality of service (QoS) information across the network.

https://t.me/learningnets

Disabling IPv4 Route Exchange

```
Router(config-router)#
no bgp default ipv4-unicast
```

- Cisco IOS and IOS XE Software only
- The exchange of IPv4 routes between BGP neighbors is enabled by default. Every configured neighbor will also receive IPv4 routes.
- This command disables the default exchange of IPv4 routes. Neighbors that need to receive IPv4 routes have to be activated for IPv4 route exchange.
- Use this command when the same router carries Internet and VPNv4 routes and you do not want to propagate Internet routes to some PE neighbors.

© 2012 Cisco and/or its affiliates. All rights reserved.

The BGP configuration that has been discussed so far is appropriate for situations where the PE routers provide Internet and VPN connectivity. If the PE routers provide only VPN connectivity, they do not need Internet routing, and the IPv4 route exchange should be disabled. Here are the two ways of disabling IPv4 route exchange:

- To disable IPv4 route exchange for only a few neighbors, your best option is to disable the IPv4 route exchange on a neighbor-by-neighbor basis by using the **no neighbor activate** command.
- To disable IPv4 route exchange for most (or all) of the neighbors, you can use the **no bgp default ipv4-unicast** command. After you enter this command, you must manually activate IPv4 route exchange for each configured global BGP neighbor.

https://t.me/learningnets

Disabling IPv4 Route Exchange (Cont.)

- Neighbor 172.16.32.14 receives only Internet routes.
- Neighbor 172.16.32.15 receives only VPNv4 routes.
- Neighbor 172.16.32.27 receives Internet and VPNv4 routes.

```
router bgp 65173
 no bgp default ipv4-unicast
 neighbor 172.16.32.14 remote-as 65173
 neighbor 172.16.32.15 remote-as 65173
 neighbor 172.16.32.27 remote-as 65173
! Activate IPv4 route exchange

address-family ipv4
 neighbor 172.16.32.14 activate
 neighbor 172.16.32.27 activate
! Step#2 - VPNv4 route exchange
 address-family vpnv4
  neighbor 172.16.32.15 activate
  neighbor 172.16.32.27 activate
```

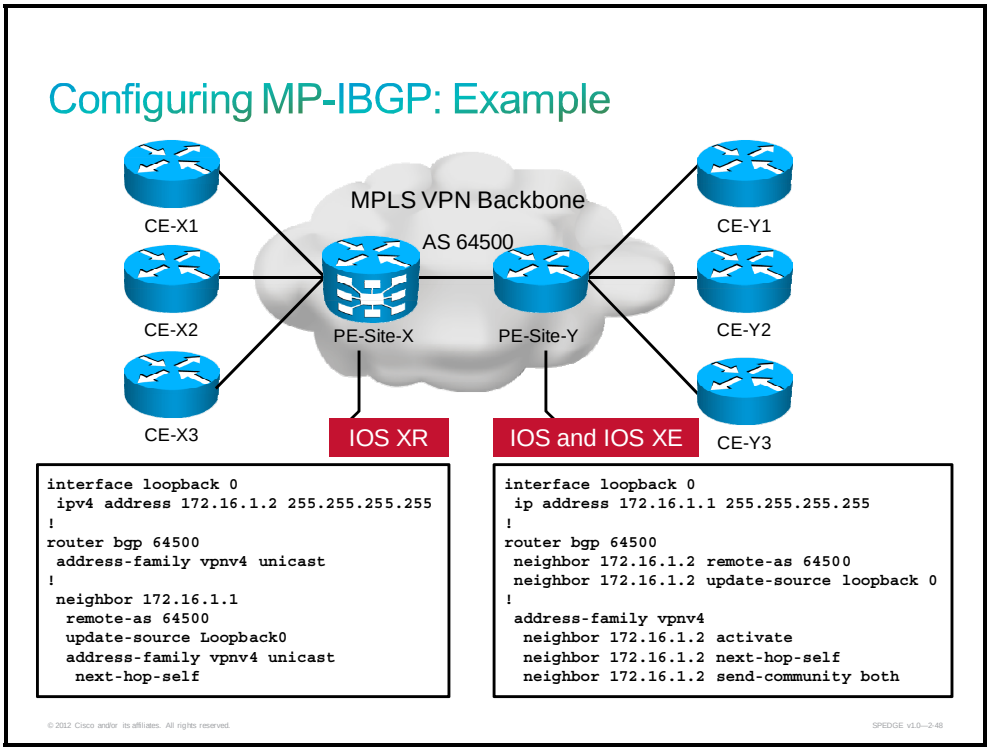
© 2012 Cisco and/or its affiliates. All rights reserved.

In this example, only a subset of BGP neighbors needs to receive IPv4 routes.

In the figure, the default propagation of IPv4 routes is therefore disabled. IPv4 route exchange—and VPNv4 route exchange—is manually activated on a neighbor-by-neighbor basis:

- Neighbor 172.16.32.14 receives only Internet routes that are based on the IPv4 activation.
- Neighbor 172.16.32.15 receives only VPNv4 routes that are based on the VPNv4 activation.
- Neighbor 172.16.32.27 receives Internet and VPNv4 routes that are based on both IPv4 and VPNv4 activations.

https://t.me/learningnets



The right box in the figure shows Cisco IOS Software configuration. A neighbor must be configured in the BGP section and then activated in the address family block. The extended community command is added automatically.

The left box shows Cisco IOS XR Software configuration. In this case, the VPNv4 address family is configured in the BGP section and then applied in the neighbor configuration block.

https://t.me/learningnets

Summary

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

Summary

- A VRF table is a routing and forwarding instance that associates additional attributes such as RD, import RT, and export RT to routing entries.
- “VRF-Lite” equals “VRF without the need to run MPLS.”
- MP-BGP is responsible for allocating labels for VPN routes and advertising them to other edge routers when using MPLS.

© 2012 Cisco and/or its affiliates. All rights reserved.

https://t.me/learningnets

https://t.me/learningnets

Lesson 2

Connecting Customers Using Simple Routing Protocols

Overview

This lesson explains provider edge (PE)-customer edge (CE) routing protocol configuration steps and the various routing protocols that you can run between PE and CE routers. These protocols include Routing Information Protocol (RIP), Enhanced Interior Gateway Routing Protocol (EIGRP), and static routes.

It is important to understand not only what you can configure between PE and CE routers when you are setting up Multiprotocol Label Switching (MPLS) VPNs, but also how to accomplish the configuration successfully. It is also very important to be able to determine what steps you should take when trying to solve a problem with your MPLS VPN network.

Objectives

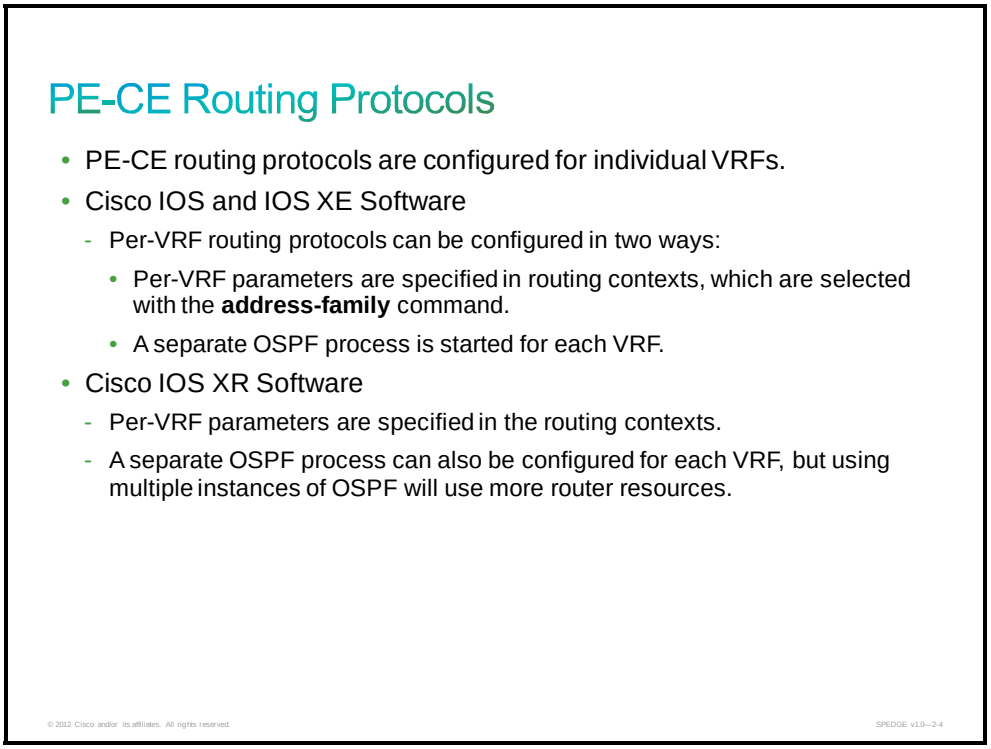
Upon completing this lesson, you will be able to describe how to configure routing protocols between PE and CE routers. You will be able to meet this objective:

- Connect customers using per-VRF static routes, RIP PE-CE routing sessions, and EIGRP PE-CE routing sessions

https://t.me/learningnets

PE-CE Routing

This topic identifies the requirements for configuring PE-CE routing protocols.



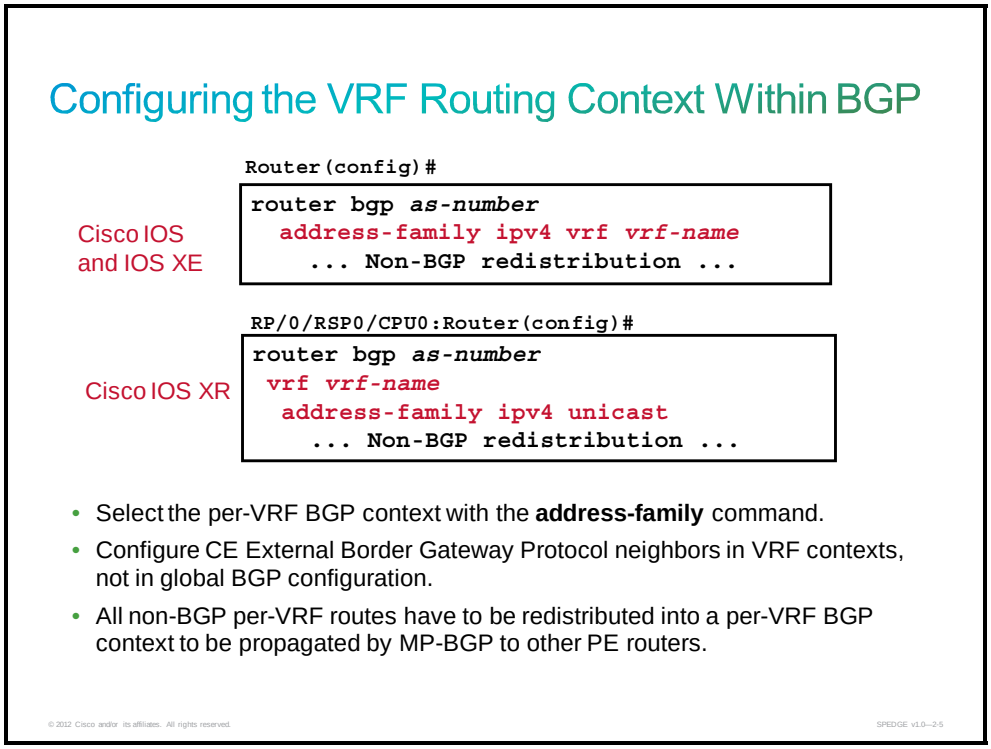
After you configure virtual routing and forwarding (VRF) instances and establish Multiprotocol Internal Border Gateway Protocol (MP-IBGP) connectivity between PE routers, you need to configure routing protocols between the PE router and the attached CE routers. The PE-CE routing protocols must be configured for individual VRFs. Sites that are in the same VPN but in different VRFs cannot share the same PE-CE routing protocol.

Note	The per-VRF configuration of the PE-CE routing protocols is another good reason for grouping as many sites into a VRF as possible.

The per-VRF routing protocols can be configured in these ways:

- Cisco IOS and IOS XE Software
  - Per-VRF routing protocols can be configured in two ways:
    - Per-VRF parameters are specified in routing contexts, which are selected with the **address-family** command.
    - A separate OSPF process has to be started for each VRF.
  - Before Cisco IOS Release 12.3(4)T, the overall number of routing processes per router was limited to 32, of which only 28 were available for VRF assignment.
- Cisco IOS XR Software
  - Per-VRF parameters are specified in the routing contexts.
  - A separate OSPF process can also be configured for each VRF, but using multiple instances of OSPF will use more router resources.

https://t.me/learningnets



On Cisco IOS and IOS XE devices, select the VRF routing context with the **address-family ipv4 vrf** *vrf-name* command in the RIP and Border Gateway Protocol (BGP) routing processes. All per-VRF routing protocol parameters (network numbers, passive interfaces, neighbors, filters, and so on) are configured under this address family.

On Cisco IOS XR devices, first define a VRF with the **vrf** *vrf-name* command in the BGP routing processes. Then select the routing context with the **address-family ipv4 unicast** command. All per-VRF routing protocol parameters (network numbers, passive interfaces, neighbors, filters, and so on) are configured under this address family.

https://t.me/learningnets

Configuring per-VRF Static Routes (Cisco IOS and IOS XE Software)

```
Router(config)#
ip route vrf vrf-name prefix mask [next-hop-address]
[interface interface-number]
```

- This command configures per-VRF static routes.
- The route is entered in the VRF table.
- You must specify a next-hop IP address if you are not using a point-to-point interface.

Sample router configuration:

```
ip route vrf Customer_ABC 10.0.0.0 255.0.0.0 10.250.0.2
!
router bgp 65173
 address-family ipv4 vrf Customer_ABC
  redistribute static
```

© 2012 Cisco and/or its affiliates. All rights reserved.

To establish static routes for a VPN VRF instance on Cisco IOS and IOS XE devices, use the **ip route vrf** command in global configuration mode. To disable static routes, use the **no** form of this command.

Configuring per-VRF Static Routes (Cisco IOS XR Software)

```
RP/0/RSP0/CPU0:Router(config)#
router static
 vrf vrf-name
  address-family ipv4 unicast
   prefix mask [next-hop-address] [interface interface-number]
```

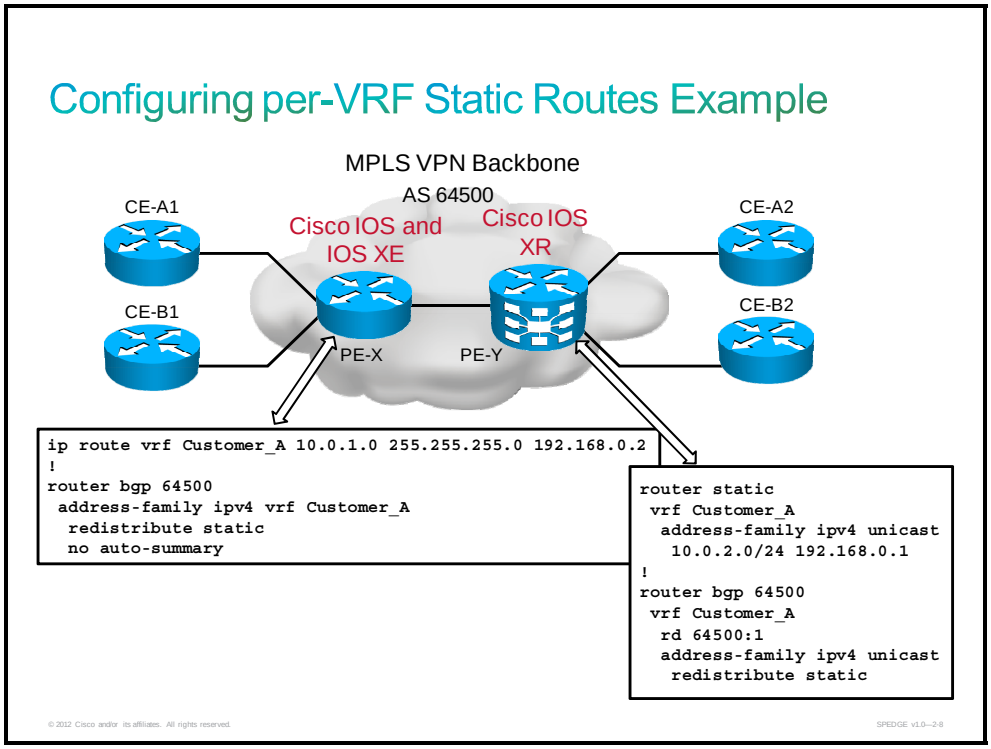
Sample router configuration:

```
router static
 vrf Customer_A
  address-family ipv4 unicast
   10.0.2.0/24 192.168.0.1
!
router bgp 64500
 vrf Customer_A
  rd 64500:1
  address-family ipv4 unicast
   redistribute static
```

© 2012 Cisco and/or its affiliates. All rights reserved.

On Cisco IOS XR devices, you must first define the static router. To enter static router configuration mode, use the **router static** command in global configuration mode. Use the **vrf** command to configure a VRF instance. Enter address family configuration mode with the **address-family ipv4 unicast** command to configure a static route.

https://t.me/learningnets



The examples in the figure (Cisco IOS and IOS XE and IOS XR CLIs) show how to configure PE-CE routing using static routes.

https://t.me/learningnets

Configuring RIP PE-CE Routing

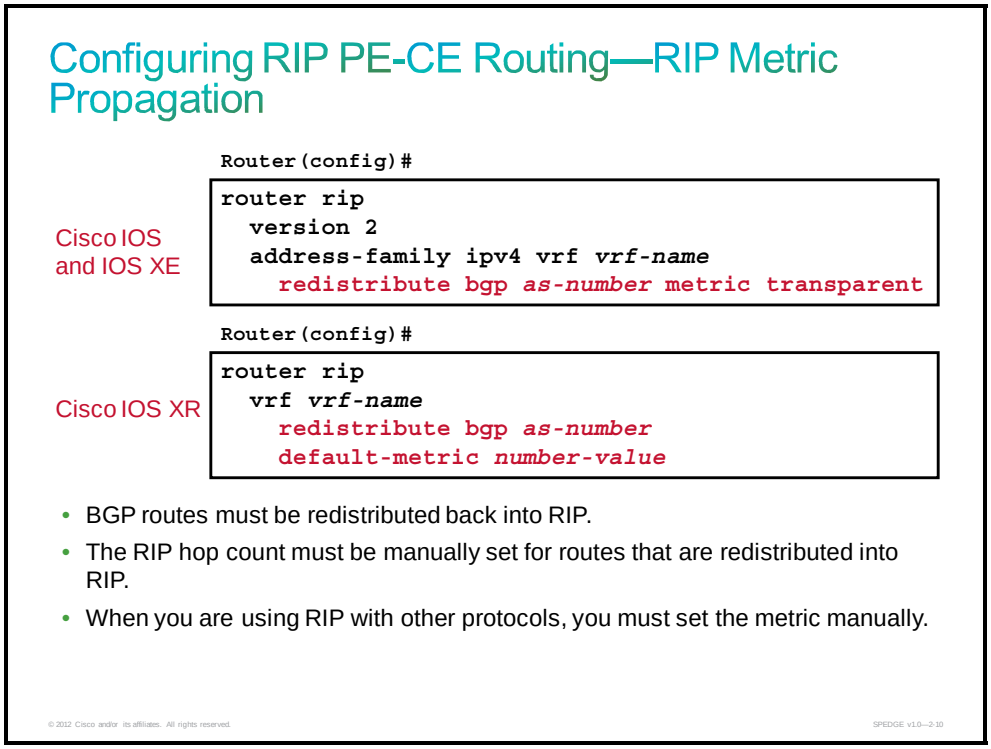
- A routing context is configured for each VRF running RIP.
- RIP parameters have to be specified in the VRF.
- Some parameters configured in the RIP process are propagated to routing contexts (for example, RIP version).
- Only RIPv2 is supported.

© 2012 Cisco and/or its affiliates. All rights reserved.

To configure RIP as the PE-CE routing protocol on Cisco IOS and IOS XE devices, start the configuration of the individual routing context with the **address-family ipv4 vrf** *vrf-name* command in router configuration mode. You can enter all standard RIP parameters in the per-VRF routing context. Global RIP parameters that are entered in the scope of RIP router configuration are inherited by each routing context and can be overwritten if needed in each routing context.

Note	Only RIP version 2 (RIPv2), and not RIP version 1 (RIPv1), is supported as the PE-CE routing protocol. It is a good practice to configure the RIP version as a global RIP parameter using the **version 2** command in router configuration mode.

https://t.me/learningnets



The interior gateway protocol (IGP) metric is always copied into the multi-exit discriminator (MED) attribute of the BGP route when an IGP route is redistributed into BGP. Within standard BGP implementation, the MED attribute is used only as a route selection criterion. The MED attribute is not copied back into the IGP metric. The IGP metric must be specified in the **redistribute** command or by using the **default-metric** command in router configuration mode.

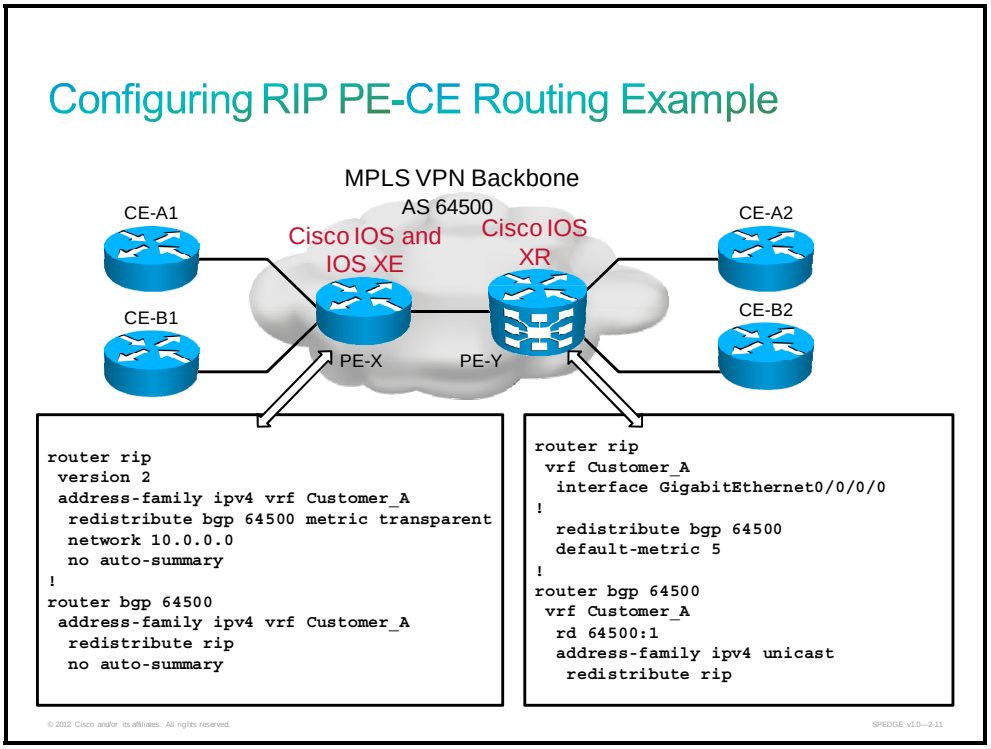
On Cisco IOS and IOS XE devices, the MPLS VPN extension to the **redistribute** command (the **metric transparent** option) allows the MED attribute to be inserted as the IGP metric of a route redistributed from BGP back into RIP. This extension gives transparent end-to-end (from the customer perspective) RIP routing:

- By default, the RIP hop count is inserted into the BGP MED attribute when the RIP route is redistributed into BGP by the ingress PE router.
- You can configure the value of the MED attribute (the original RIP hop count) to be copied into the RIP hop count when the BGP route is redistributed back into RIP. This action causes the whole MPLS VPN backbone to appear as a single hop to the CE routers.

Note	You should not change the MED value within BGP if you use the **redistribute metric transparent** command.

On Cisco IOS XR devices, use the **default-metric** command in VRF configuration mode for RIP routing. The default metric value (*number-value*) is a number from 1 to 15.

https://t.me/learningnets



The examples in the figure (Cisco IOS, IOS XE, and IOS XR CLIs) show how to configure PE-CE routing using RIP as a routing protocol.

https://t.me/learningnets

Configuring EIGRP PE-CE Routing

- Provides EIGRP with the capability to redistribute routes through a VPN cloud.
- EIGRP extended community attributes are used to define EIGRP routes and preserve internal metrics.
- Supports SOO capabilities to filter VPN traffic.

© 2012 Cisco and/or its affiliates. All rights reserved.

The MPLS VPN Support for EIGRP between Provider Edge and Customer Edge feature provides the capability to transparently connect EIGRP customer networks through an MPLS-enabled BGP core network. EIGRP routes can then be redistributed through the VPN across the BGP network as internal BGP (IBGP) routes. The configuration of this feature does not require any customer equipment upgrades or configuration changes. This feature is configured only on PE routers within the service provider network.

Customer networks and remote sites are connected to each other through the MPLS VPN. The configuration of this feature allows several EIGRP sites to connect seamlessly and appear as a single network. This integration is transparent to the customer sites. When this feature is enabled, EIGRP routes are converted to IBGP routes and transported through the BGP core network. EIGRP extended community attributes are used to define EIGRP routes and preserve internal metrics. These attributes are carried across the core network by multiprotocol BGP (MP-BGP).

This feature also introduces EIGRP support for MPLS and BGP extended community attributes such as Site of Origin (SOO).

https://t.me/learningnets

Configuring EIGRP PE-CE Routing—EIGRP Metric Propagation

Cisco IOS and IOS XE

```
Router(config)#
router eigrp autonomous-system-number
address-family ipv4 vrf vrf-name
   autonomous-system as-number
   redistribute bgp as-number metric metric-value
```

Cisco IOS XR

```
Router(config)#
router eigrp autonomous-system-number
 vrf vrf-name
  address-family ipv4
   autonomous-system as-number
   redistribute bgp as-number metric metric-value
```

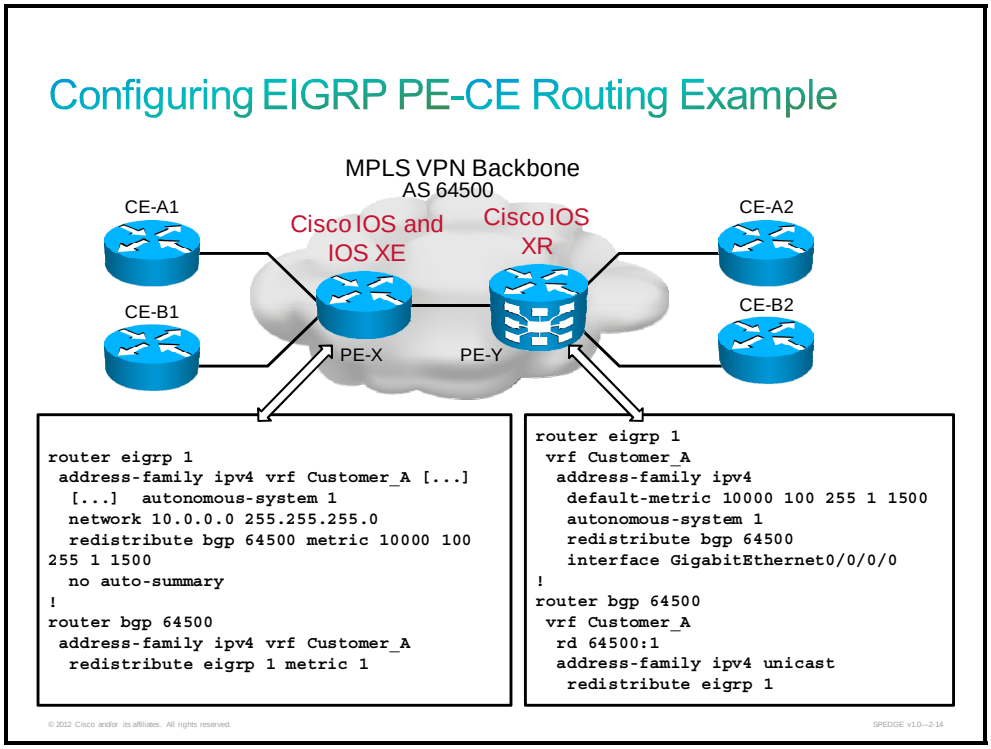
- Enables the EIGRP AS number of the CE under the address family.
- Configures per-instance AS number.
- Configures router redistribution.
- External routes that are received without the configured metric are not to be advertised to the CE router.

© 2012 Cisco and/or its affiliates. All rights reserved.

The IGP metric is always copied into the MED attribute of the BGP route when an IGP route is redistributed into BGP. Within a standard BGP implementation, the MED attribute is used only as a route-selection criterion. The MED attribute is not copied back into the IGP metric. The metric must be configured for routes from external EIGRP autonomous systems and non-EIGRP networks before these routes can be redistributed into an EIGRP CE router. The metric can be configured in the redistribute statement using the **redistribute** (IP) command or configured with the **default-metric** (EIGRP) command.

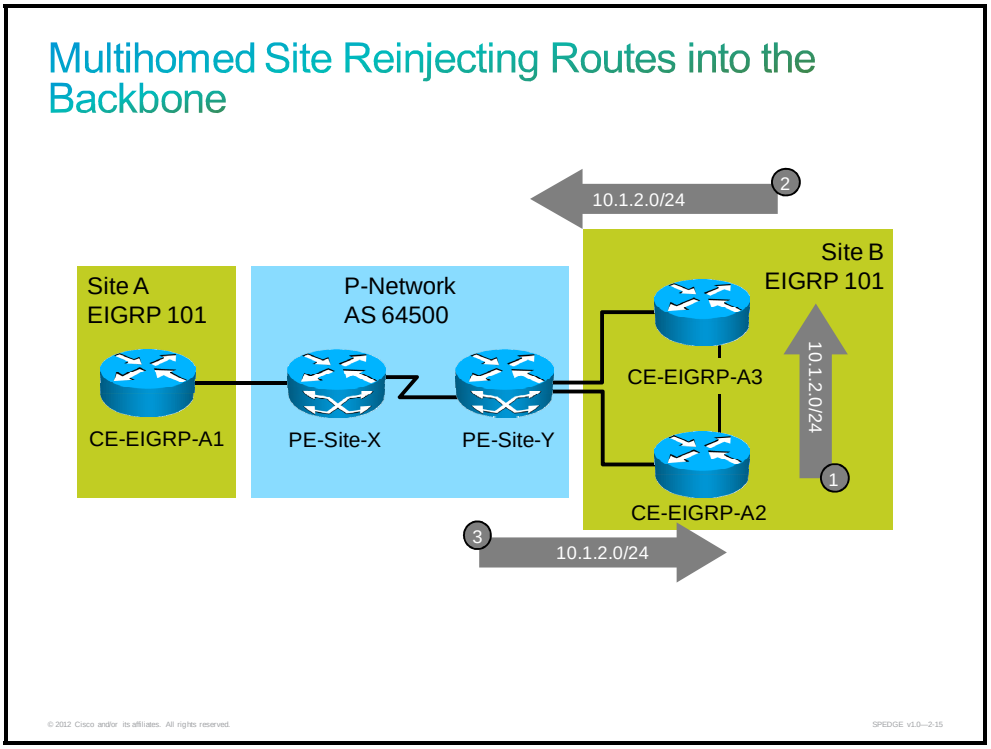
Note In an MPLS VPN environment, the original EIGRP metrics must be carried inside MP-BGP updates. This configuration is achieved by using BGP extended community attributes to carry and preserve EIGRP metrics when crossing the MP-IBGP domain. These communities define the intrinsic characteristics that are associated with EIGRP, such as the AS number or EIGRP cost metric (bandwidth, delay, load, reliability, and MTU, for example).

https://t.me/learningnets



The examples in the figure (Cisco IOS, IOS XE, and IOS XR CLIs) show how to configure PE-CE routing using EIGRP as the routing protocol.

https://t.me/learningnets



Routing loops and suboptimal routing generally occur because of mutual redistribution taking place between EIGRP PE-CE and MP-BGP in an MPLS VPN environment. Routing loops can occur in the following scenarios:

- A route that is received by a multihomed site from the backbone through one link can be forwarded back to the backbone through the other link.
- A route that originated in a multihomed site and that was sent to the backbone through one link can come back through the other link.

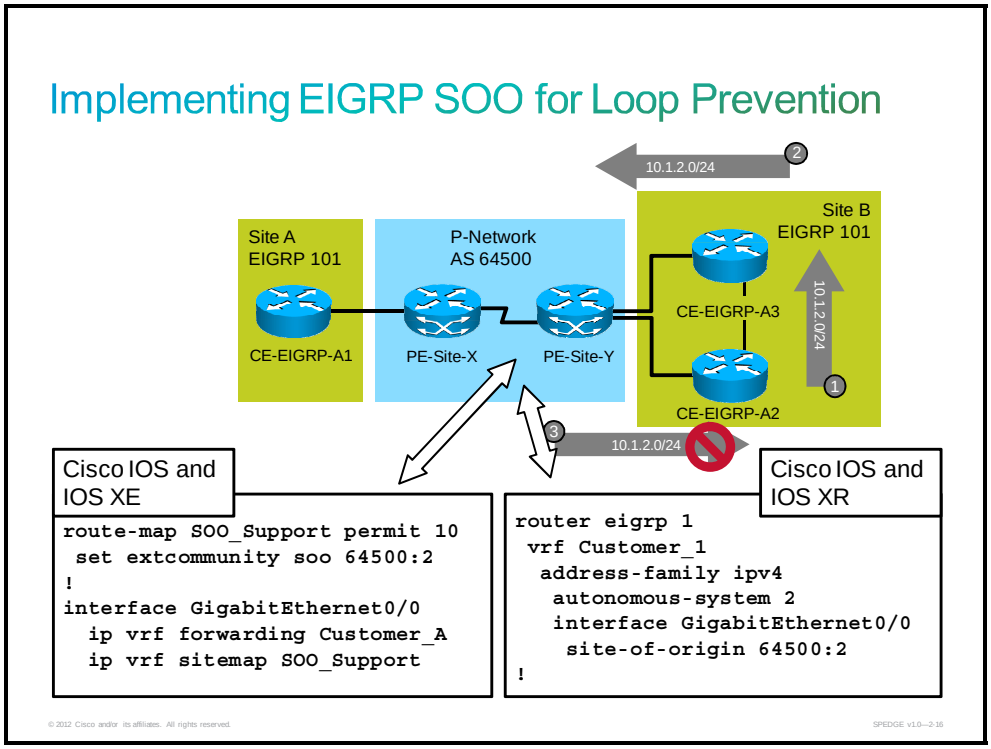
The figure shows an MPLS VPN network for a customer that has two sites, Site A and Site B. Site B is multihomed. The figure shows that EIGRP route 10.1.2.0/24 received by the multihomed site (Site B) is redistributed into the backbone at PE-Site-Y.

Routing loops and suboptimal routing can be avoided by using the following:

- The BGP Cost Community feature, which can be used to force BGP to compare locally originated EIGRP routes and MP-IBGP routes, based on the EIGRP metric
- The EIGRP SOO feature on PE and CE routers, which can be used to prevent routing loops

Note	The SOO attribute is needed only for customer networks with multihomed sites. Loops can never occur in customer networks that have only stub sites.

https://t.me/learningnets



The configuration of the SOO extended community allows routers that support this feature to identify the site from which each route originated. When this feature is enabled, the EIGRP routing process on the PE or CE router checks each received route for the SOO extended community and filters based on the following conditions:

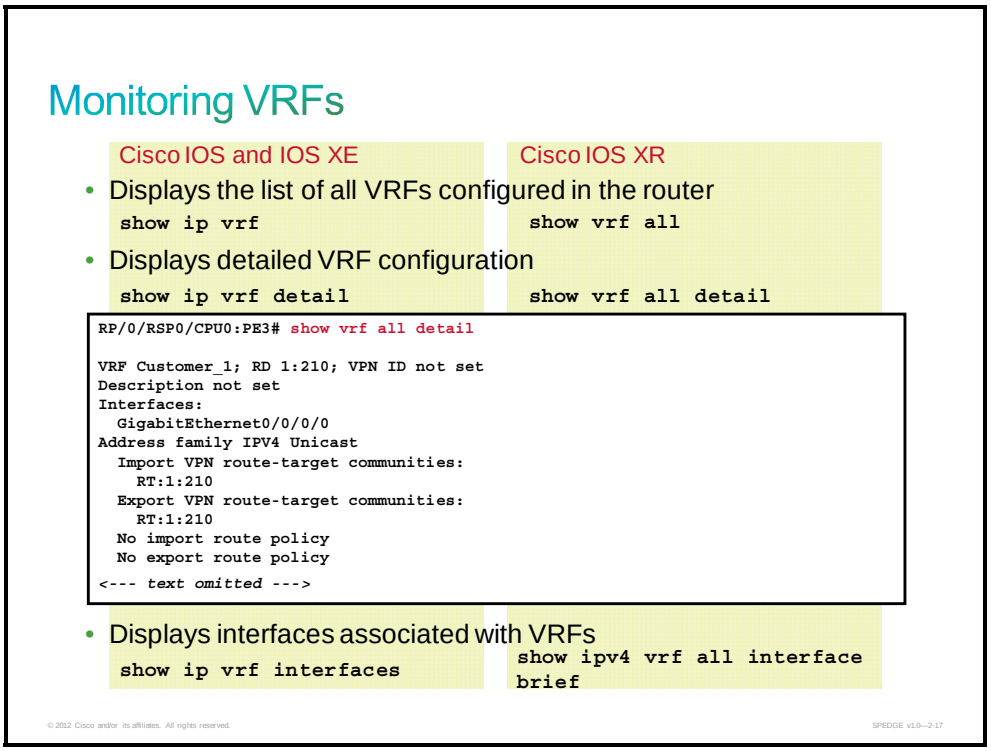
- A received route from BGP or a CE router contains an SOO value that matches the SOO value on the receiving interface:
  - If a route is received with an associated SOO value that matches the SOO value that is configured on the receiving interface, the route is filtered out because it was learned from another PE router or from a backdoor link. This behavior is designed to prevent routing loops.
- A received route from a CE router is configured with an SOO value that does not match:
  - If a route is received with an associated SOO value that does not match the SOO value that is configured on the receiving interface, the route is accepted into the EIGRP topology table so that it can be redistributed into BGP.
  - If the route is already installed in the EIGRP topology table but is associated with a different SOO value, the SOO value from the topology table is used when the route is redistributed into BGP.
- A received route from a CE router does not contain an SOO value:
  - If a route is received without an SOO value, the route is accepted into the EIGRP topology table. The SOO value from the interface that is used to reach the next-hop CE router is appended to the route before it is redistributed into BGP.

When BGP and EIGRP peers that support the SOO extended community receive these routes, they also receive the associated SOO values and pass them to other BGP and EIGRP peers that support the SOO extended community. This filtering is designed to prevent transient routes from being relearned from the originating site, which prevents transient routing loops from occurring.

https://t.me/learningnets

In conjunction with BGP Cost Community, EIGRP, BGP, and the routing information base (RIB) ensure that paths over the MPLS VPN core are preferred over backdoor links.

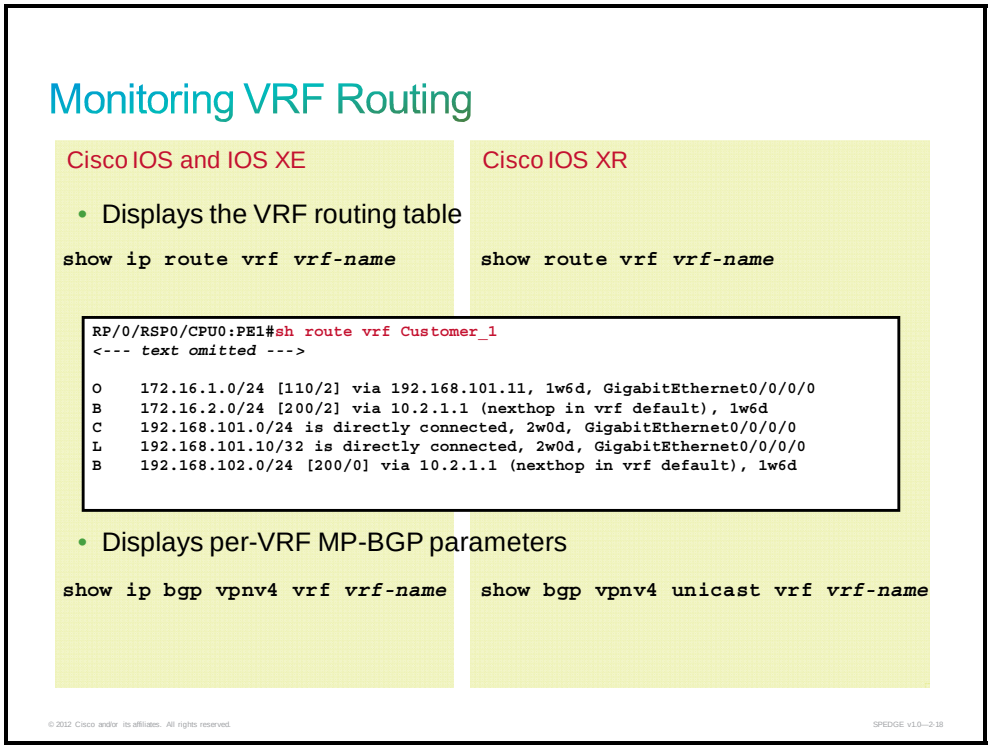
In the example in the figure, a route map, SOO_Support, with a specific SOO value of 64500:2 is defined. The newly defined route map is then applied to VRF Customer-EIGRP-A1, which connects the EIGRP neighbor (the CE-EIGRP-A2 router) to the PE-Site-Y router.



To display the set of defined VRFs, use the **show ip vrf** (IOS and IOS XE) or **show vrf all** (IOS XR) command in EXEC mode.

To display the interfaces that are associated with a specific VRF, use the **show ip vrf interfaces** (IOS and IOS XE) or **show ipv4 vrf all interface brief** (IOS XR) command.

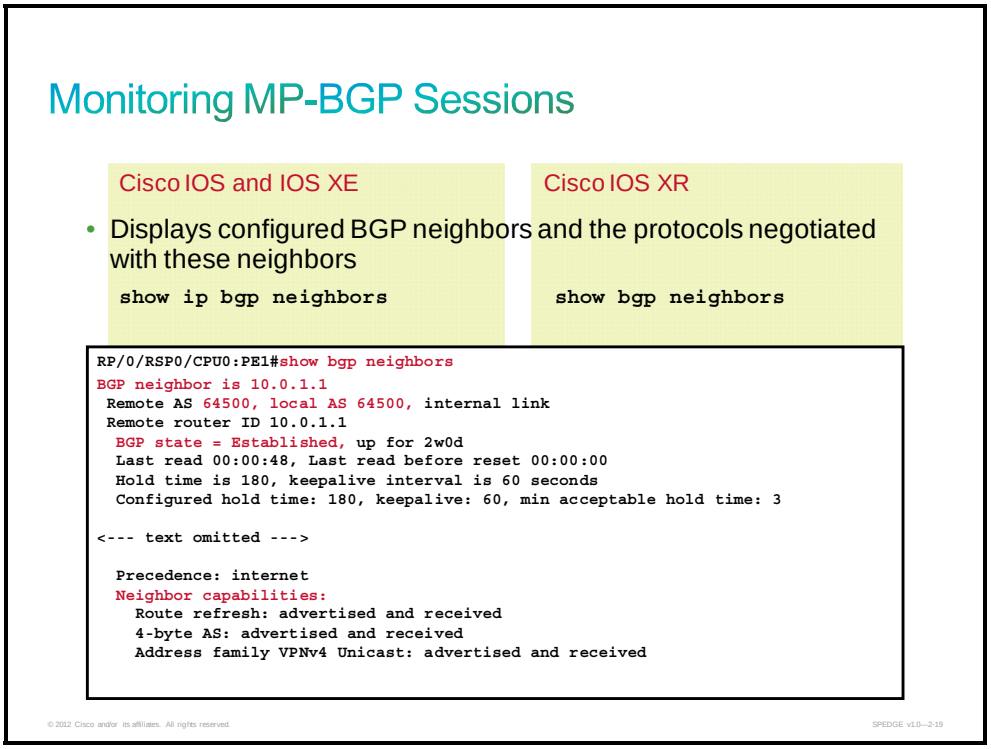
https://t.me/learningnets



These two commands can be used to monitor VRF routing:

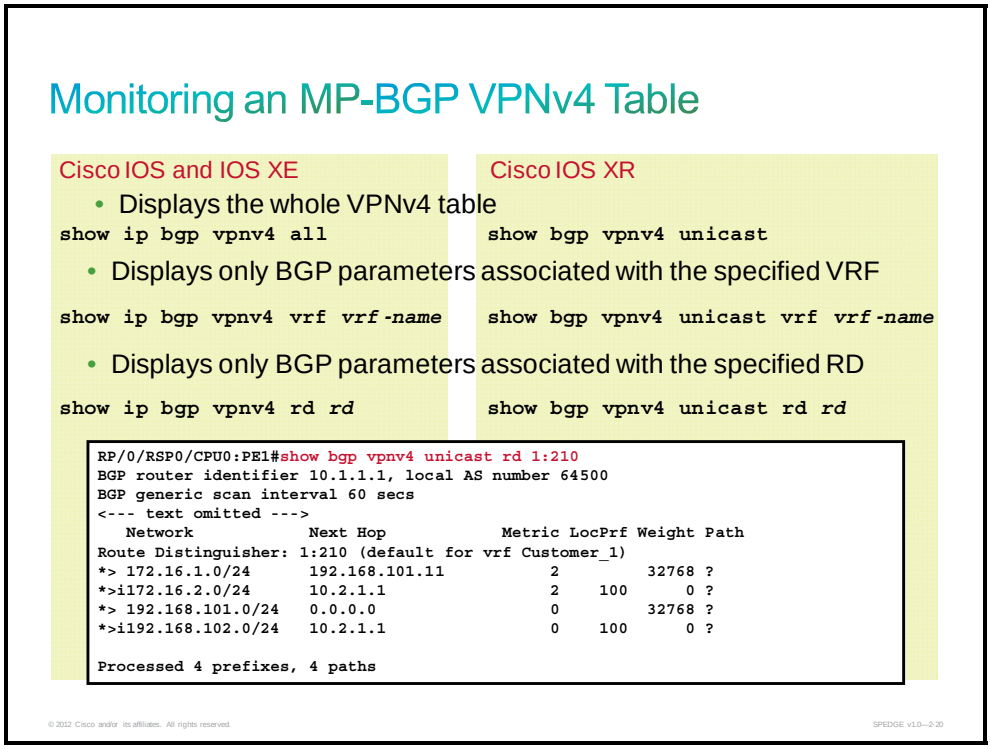
- The **show ip route vrf** (IOS and IOS XE) or **show route vrf** (IOS XR) command displays the VRF routing table.
- The **show ip bgp vpnv4 vrf** (IOS and IOS XE) or **show bgp vpnv4 unicast vrf** (IOS XR) command displays the VRF BGP table.

https://t.me/learningnets



The **show ip bgp neighbors** (Cisco IOS and IOS XE) or **show bgp neighbors** (Cisco IOS XR) command is described in detail in the Cisco IOS and Cisco IOS XR Software documentation. This command is used to monitor BGP sessions with other PE routers and the address families that are negotiated with these neighbors.

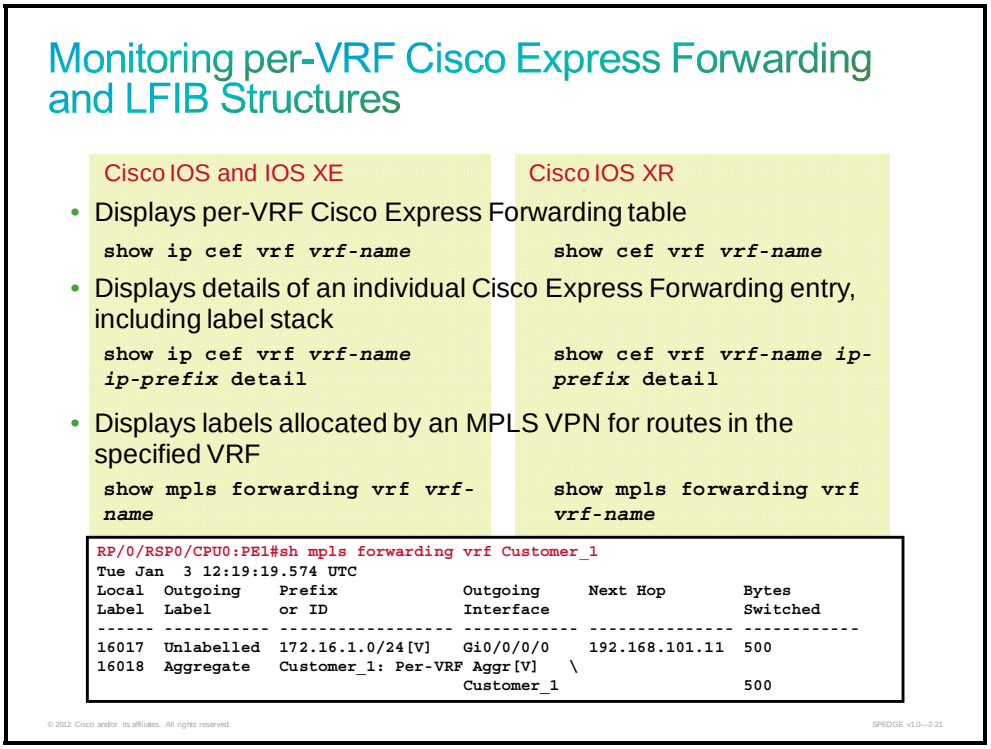
https://t.me/learningnets



The **show ip bgp vpnv4** (Cisco IOS and IOS XE) or **show bgp vpnv4 unicast** (Cisco IOS XR) command displays IPv4 BGP information and VPNv4 BGP information. To display VPNv4 BGP information on devices that are running Cisco IOS or IOS XE, use one of these keywords:

- **all** to display the whole contents of the VPNv4 BGP table
- **vrf** *vrf-name* to display VPNv4 information that is associated with the specified VRF
- **rd** *route-distinguisher* to display VPNv4 information that is associated with the specified RD

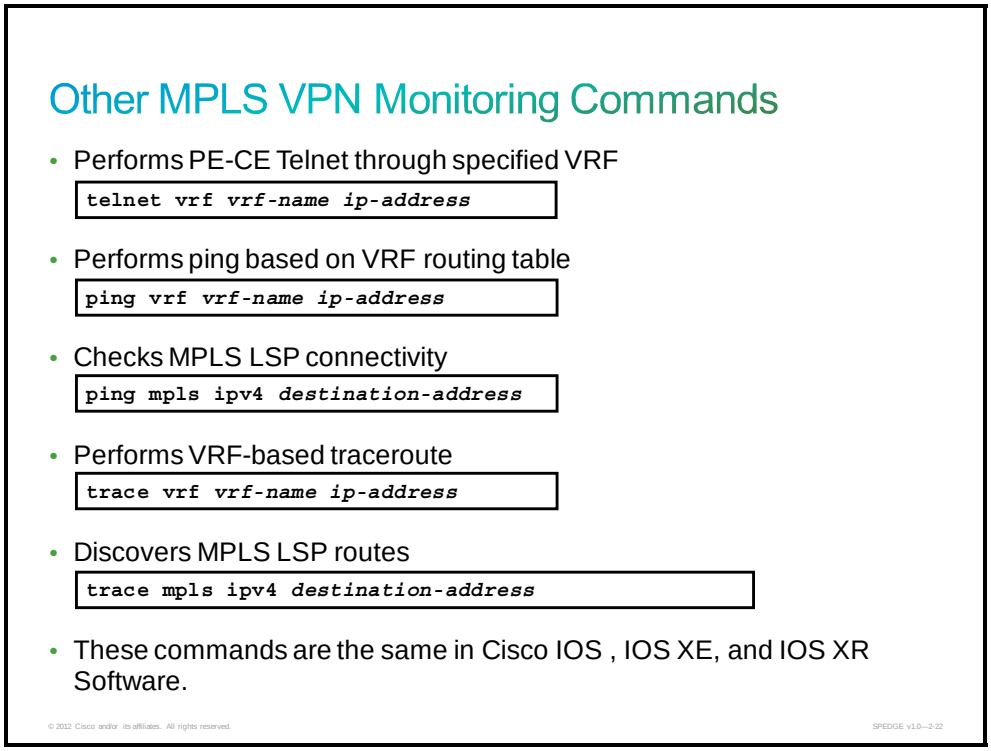
https://t.me/learningnets



These three commands can be used to display per-VRF forwarding information base (FIB) and label forwarding information base (LFIB) structures:

- The **show ip cef vrf** (IOS and IOS XE) or **show cef vrf** (IOS XR) command displays the VRF FIB.
- The **show ip cef vrf detail** (IOS and IOS XE) or **show cef vrf detail** (IOS XR) command displays detailed information about a single entry in the VRF FIB.
- The **show mpls forwarding vrf** (Cisco IOS, IOS XE, and IOS XR) command displays all labels that are allocated to VPN routes in the specified VRF.

https://t.me/learningnets



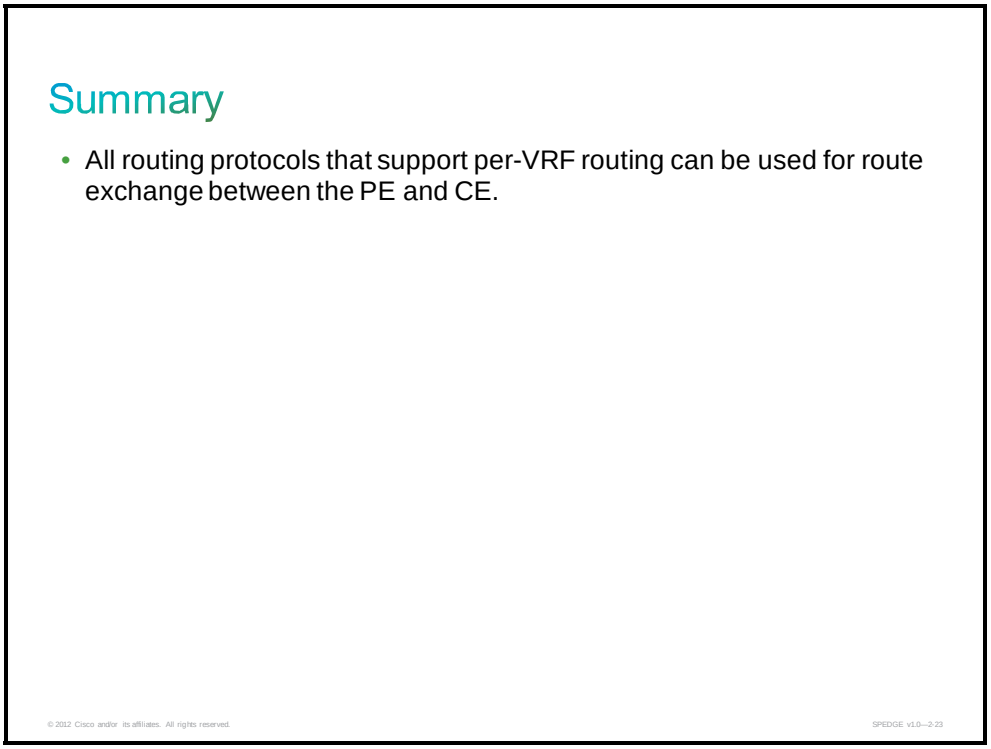
Additional Cisco IOS Software monitoring commands are as follows:

- The **telnet** command can be used to connect to a CE router from a PE router using the **/vrf** option.
- The **ping vrf** command can be used to ping a destination host reachable through a VRF.
- The **ping mpls** command can be used to check MPLS label-switched path (LSP) connectivity.
- The **trace vrf** command can be used to trace a path toward a destination reachable through a VRF.
- The **trace mpls** command can be used to discover the MPLS LSP routes that packets actually take when traveling to their destinations.

https://t.me/learningnets

Summary

This topic summarizes the primary point that was discussed in this lesson.



https://t.me/learningnets

Lesson 3

Connecting Customers Using BGP or OSPF

Overview

This lesson explains the provider edge (PE)-customer edge (CE) advanced routing protocol configuration steps and the various routing protocols that you can run between PE and CE routers.

The Open Shortest Path First (OSPF) and Border Gateway Protocol (BGP) routing protocols are discussed. This lesson describes the steps that are needed for Multiprotocol Label Switching (MPLS) VPN troubleshooting.

Objectives

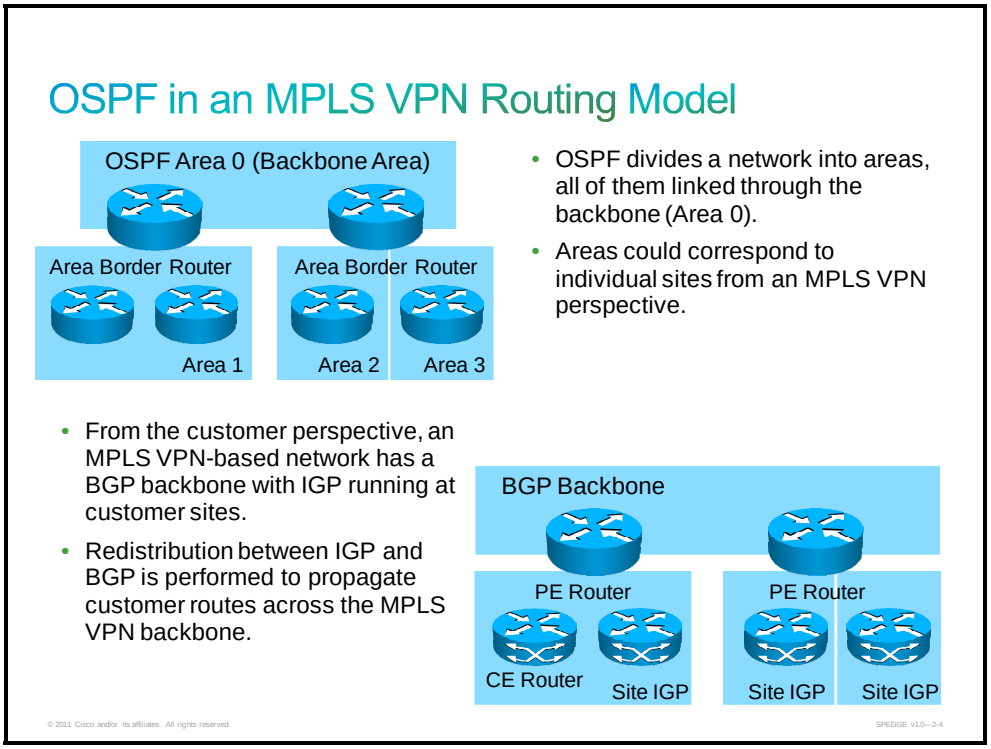
Upon completing this lesson, you will be able to describe how to configure and troubleshoot routing protocols between PE and CE routers. You will be able to meet these objectives:

- Configure an OSPF PE-CE routing session
- Configure a BGP PE-CE routing session
- Describe how to troubleshoot MPLS VPNs

https://t.me/learningnets

OSPF as the PE-CE Routing Protocol

This topic describes OSPF as the routing protocol between PE and CE routers.



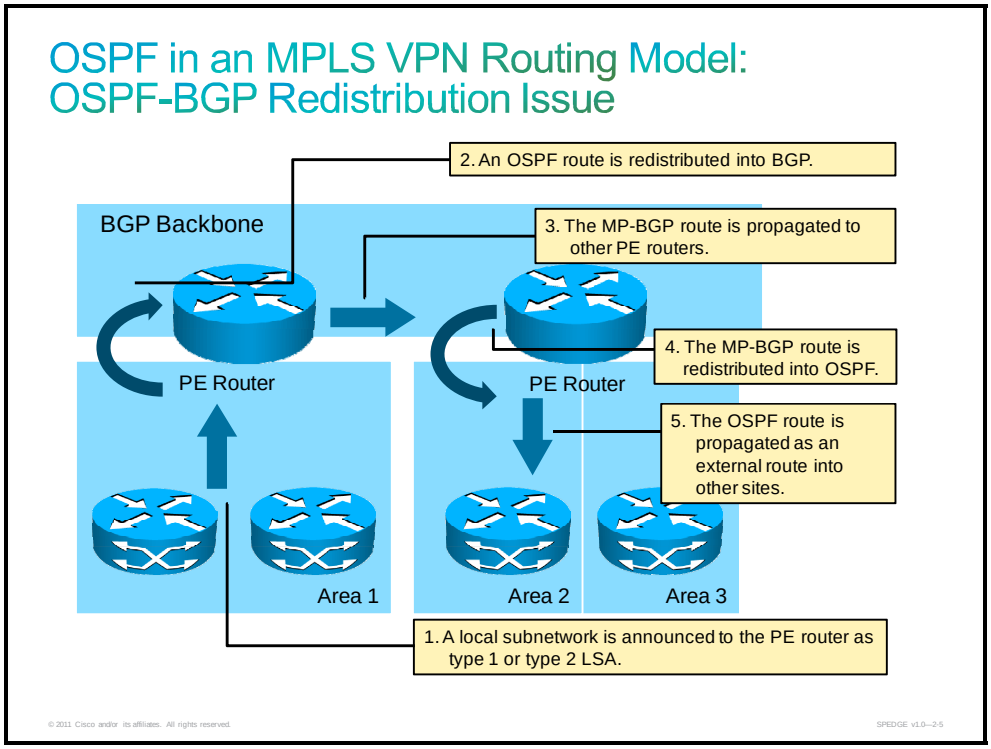
The OSPF routing protocol was designed to support hierarchical networks with a central backbone. A network that is running OSPF is divided into areas. All areas must be directly connected to the backbone area (Area 0). The whole OSPF network (backbone area and other connected areas) is called the *OSPF domain*.

The OSPF areas in the customer network can correspond to individual sites, but these other options are often encountered:

- A single area could span multiple sites (for example, the customer decides to use one area per region, but the region contains multiple sites).
- The backbone area could be extended into individual sites.

The MPLS VPN routing model introduces a BGP backbone into the customer network. Isolated copies of the interior gateway protocol (IGP) run at every site, and Multiprotocol BGP (MP-BGP) is used to propagate routes between sites. Redistribution between the customer IGP (running between PE routers and CE routers) and the backbone MP-BGP is performed at every PE router.

https://t.me/learningnets



The IGP-BGP redistribution that is introduced by the MPLS VPN routing model does not fit well into customer networks running OSPF. When an OSPF customer is migrated to an MPLS VPN implementation, any route that is redistributed into OSPF from another routing protocol will now be redistributed as an *external* OSPF route. The OSPF routes received by one PE router are propagated across the MPLS backbone and redistributed back into OSPF at another site as external OSPF routes.

Note Remember that link-state advertisement (LSA) 1 and LSA 2 never leave the OSPF area.

https://t.me/learningnets

OSPF in an MPLS VPN Routing Model: Caveats

- The OSPF route type is not preserved when the OSPF route is redistributed into BGP.
- All OSPF routes from a site are inserted as external (type 5 LSA) routes into other sites.
- The result is that OSPF route summarization and stub areas are hard to implement.

Conclusion: MPLS VPNs must extend the classic OSPF-BGP routing model.

© 2011 Cisco and/or its affiliates. All rights reserved.

With the traditional OSPF-BGP redistribution, the OSPF route type (internal or external route) is not preserved when the OSPF route is redistributed into BGP. When that route is redistributed back into OSPF, it is always redistributed as an external OSPF route.

This list identifies some of the caveats that are associated with external OSPF routes:

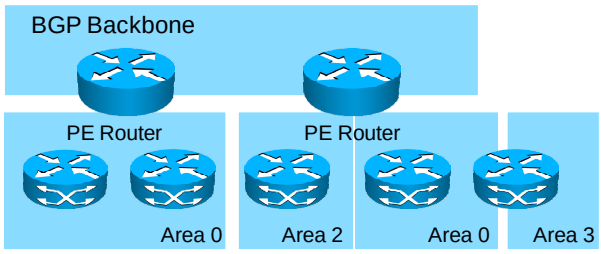
- External routes cannot be summarized.
- External routes are flooded across all OSPF areas.
- External routes could use a different metric type that is not comparable to OSPF cost.
- External routes are not inserted in stub areas or not-so-stubby areas (NSSAs).
- Internal routes are always preferred over external routes, regardless of their cost.

Because of all these caveats, migrating an OSPF customer toward an MPLS VPN service might have a severe impact on the routing of that customer. The MPLS VPN architecture must therefore extend the classic OSPF-BGP routing model to support transparent customer migration.

https://t.me/learningnets

OSPF Superbackbone

- OSPF Area 0 might extend into individual sites.
- The MPLS VPN backbone has to become a superbackbone for OSPF.



- OSPF between sites will not use normal OSPF-BGP redistribution.
- OSPF continuity must be provided across the MPLS VPN backbone:
  - Internal OSPF routes should remain internal OSPF routes.
  - External routes should remain external routes.
  - OSPF metrics should be preserved.
- CE routers run standard OSPF software.

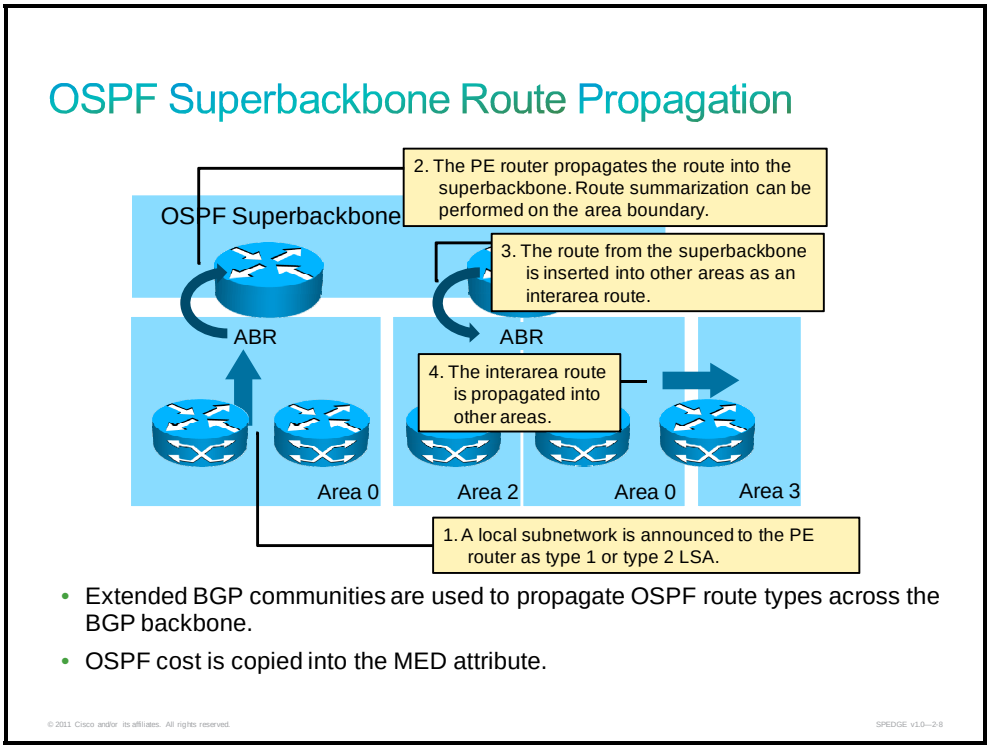
© 2011 Cisco and/or its affiliates. All rights reserved.

The MPLS VPN architecture extends the OSPF architecture by introducing another backbone above OSPF Area 0, the superbackbone. The OSPF superbackbone is implemented with MP-BGP between the PE routers but is otherwise transparent to the OSPF routers. The architecture even allows disjointed OSPF backbone areas (Area 0) at MPLS VPN customer sites.

These goals must be met by the OSPF superbackbone:

- The superbackbone will not use standard OSPF-BGP redistribution.
- OSPF continuity must be provided between OSPF sites:
  - Internal OSPF routes must remain internal OSPF routes.
  - External OSPF routes must remain external OSPF routes.
  - Non-OSPF routes that are redistributed into OSPF must appear as external OSPF routes in OSPF.
  - OSPF metrics and metric types (external 1 or external 2) must be preserved.
- The OSPF superbackbone will be transparent to the CE routers that run standard OSPF software.

https://t.me/learningnets



The MPLS VPN superbackbone appears as another layer of hierarchy in the OSPF architecture. The PE routers that connect regular OSPF areas to the superbackbone therefore appear as OSPF Area Border Routers (ABRs) in the OSPF areas to which they are attached. ABRs also appear as Autonomous System Boundary Routers (ASBRs) in nonstub areas.

From the perspective of a standard OSPF-speaking CE router, the PE routers insert interarea routes from other areas into the area in which the CE router is present. The CE routers are not aware of the superbackbone or of other OSPF areas present beyond the MPLS VPN superbackbone.

With the OSPF superbackbone architecture, the continuity of OSPF routing is preserved:

- The OSPF intra-area route, described in the OSPF router LSA or network LSA, is inserted into the OSPF superbackbone by redistributing the OSPF route into MP-BGP. Route summarization can be performed on the redistribution boundary by the PE router.
- The MP-BGP route is propagated to other PE routers and inserted as an OSPF route into other OSPF areas.

The OSPF superbackbone is implemented with the help of several BGP attributes:

- A new BGP extended community was defined to carry the OSPF route type and OSPF area across the BGP backbone.
- As in the standard OSPF-BGP redistribution, the OSPF cost is carried in the multi-exit discriminator (MED) attribute.

https://t.me/learningnets

OSPF Superbackbone Implementation

Internal OSPF routes

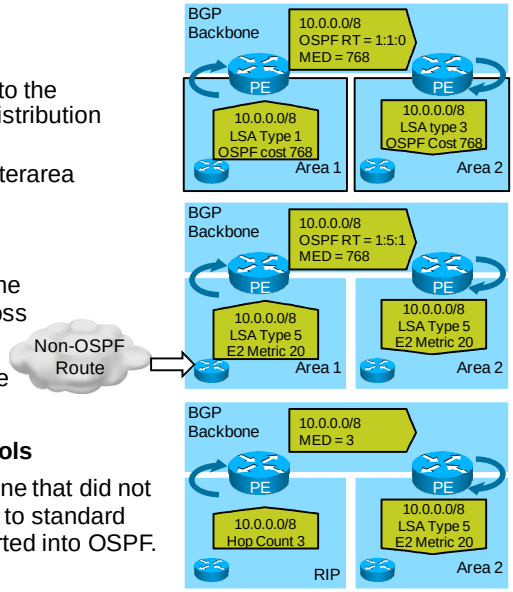
- The OSPF route type is copied into the extended BGP community on redistribution into BGP.
- The egress PE router performs interarea transformation.

External OSPF routes

- Routes are propagated in the same way as internal OSPF routes across the superbackbone.
- The external metric and route type are preserved.

Routes from other routing protocols

- Routes from the MP-BGP backbone that did not originate in OSPF are still subject to standard redistribution behavior when inserted into OSPF.



© 2011 Cisco and/or its affiliates. All rights reserved.

The first diagram in the figure illustrates the propagation of internal OSPF routes across the MPLS VPN superbackbone.

In the second diagram, the external OSPF routes are redistributed into MP-BGP in the same way as the internal OSPF routes.

In the third diagram, the MPLS VPN superbackbone retains the traditional OSPF-BGP route redistribution behavior for routes that did not originate in OSPF at other sites.

https://t.me/learningnets

Configuring PE-CE OSPF Routing

Follow these steps to configure OSPF as the PE-CE routing protocol:

- Configure a per-VRF copy of OSPF.
- Configure redistribution of MP-BGP into OSPF.
- Configure redistribution of OSPF into MP-BGP.

© 2011 Cisco and/or its affiliates. All rights reserved.

To configure OSPF as a PE-CE routing protocol, you must start a separate OSPF process for each virtual routing and forwarding (VRF) instance in which you want to run OSPF. The per-VRF OSPF process is configured in the same way as a standard OSPF process. You can use all the OSPF features that are available in Cisco IOS Software.

You need to redistribute OSPF routes into BGP and redistribute BGP routes into OSPF if necessary. Alternatively, you can originate a default route into a per-VRF OSPF process by using the **default-information originate always** command in router configuration mode.

MP-BGP propagates more than just OSPF cost across the MPLS VPN backbone. The propagation of additional OSPF attributes into MP-BGP is automatic and requires no extra configuration.

https://t.me/learningnets

Configuring PE-CE OSPF Routing: Cisco IOS and IOS XE Software

```
router(config)#
router ospf process-id vrf vrf-name
  ... Standard OSPF parameters ...
```

- This command starts the per-VRF OSPF routing process.

```
router(config-router)#
redistribute bgp as-number subnets
```

- This command redistributes MP-BGP routes into OSPF. The **subnets** keyword is mandatory for proper operation.

```
router(config)#
router bgp as-number
  address-family ipv4 vrf vrf-name
    redistribute ospf process-id [match [internal]
[external-1] [external-2]]
```

- OSPF-BGP route redistribution is configured with the **redistribute** command under the proper **address-family** command.

© 2011 Cisco and/or its affiliates. All rights reserved.

In an MPLS VPN deployment, each VPN VRF needs a separate OSPF process when configured to run OSPF. OSPF support for unlimited software VRFs per PE router addresses the scalability issue for OSPF VPNs by eliminating the OSPF VPN limit of 32 processes.

To configure an OSPF routing process within a VRF on Cisco IOS and IOS XE devices, use the **router ospf** command in global configuration mode. OSPF-BGP route redistribution is configured with the **redistribute** command under the proper **address-family** command. Without the OSPF **match** keyword specified, only internal OSPF routes are redistributed into BGP.

https://t.me/learningnets

Configuring PE-CE OSPF Routing: Cisco IOS XR Software

```
RP/0/RP0/CPU0:router(config-ospf)#
vrf vrf-name
  ... Standard OSPF parameters ...
```

- This command starts the per-VRF OSPF routing process.

```
RP/0/RP0/CPU0:router(config-ospf-vrf)#
redistribute bgp as-number
```

- This command redistributes MP-BGP routes into OSPF.

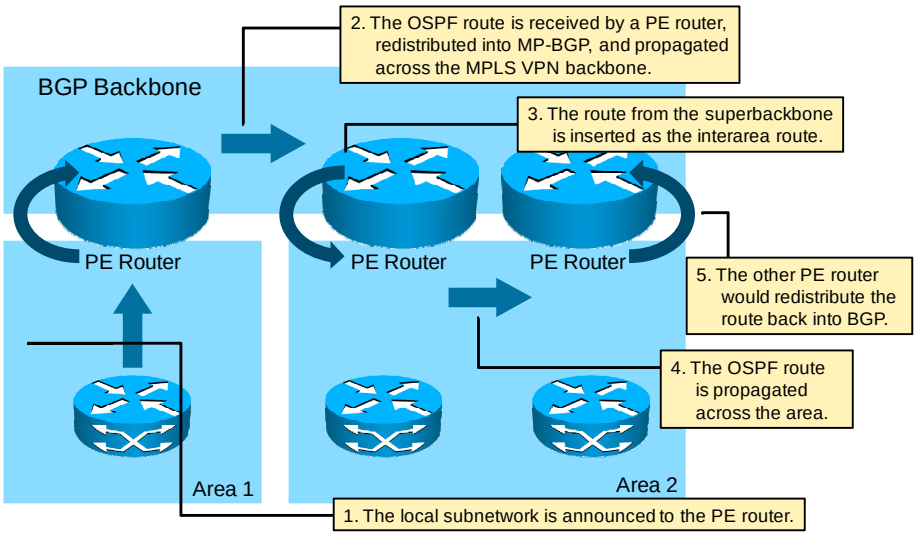
```
RP/0/RP0/CPU0:router(config)#
router bgp as-number
 vrf vrf-name
  address-family ipv4 unicast
   redistribute ospf process-id [match {external [1|2] |
internal}]
```

- OSPF-BGP route redistribution is configured under the proper **address-family** command.

© 2011 Cisco and/or its affiliates. All rights reserved.

To configure an OSPF routing process within a VRF on Cisco IOS XR devices, you must first enter VRF configuration mode for OSPF routing. OSPF-BGP route redistribution is configured with the **redistribute** command under the proper **address-family** command.

Routing Loops Between MP-BGP and OSPF



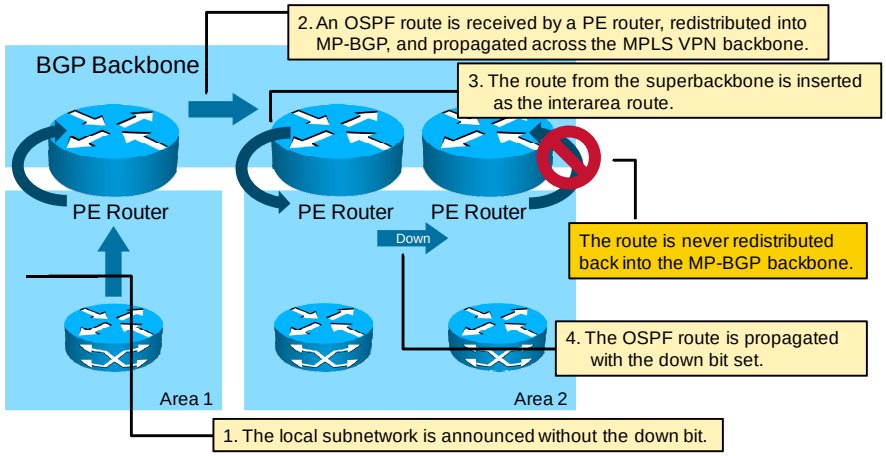
© 2011 Cisco and/or its affiliates. All rights reserved.

OSPF developers took many precautions to avoid routing loops between OSPF areas. For example, intra-area routes are always preferred over interarea routes. These rules do not work when the superbackbone is introduced. For example, consider the network in the figure, where the receiving OSPF area has two PE routers attached to it.

https://t.me/learningnets

OSPF Down Bit

- A down bit has been introduced in the options field of the OSPF LSA header.
- PE routers set the down bit when redistributing routes from MP-BGP into OSPF.
- PE routers never redistribute OSPF routes with the down bit set into MP-BGP.

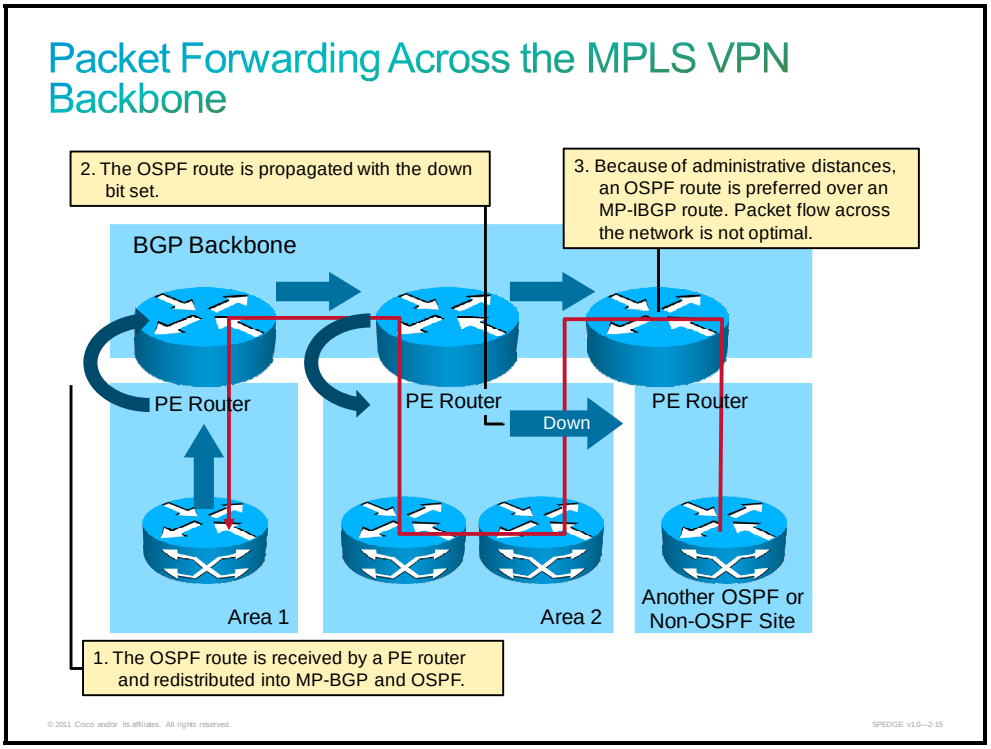


© 2011 Cisco and/or its affiliates. All rights reserved.

The down bit was introduced as a mechanism to prevent route redistribution loops between OSPF and MP-BGP that is running between PE routers. The OSPF down bit is part of the Options field in the OSPF LSA header.

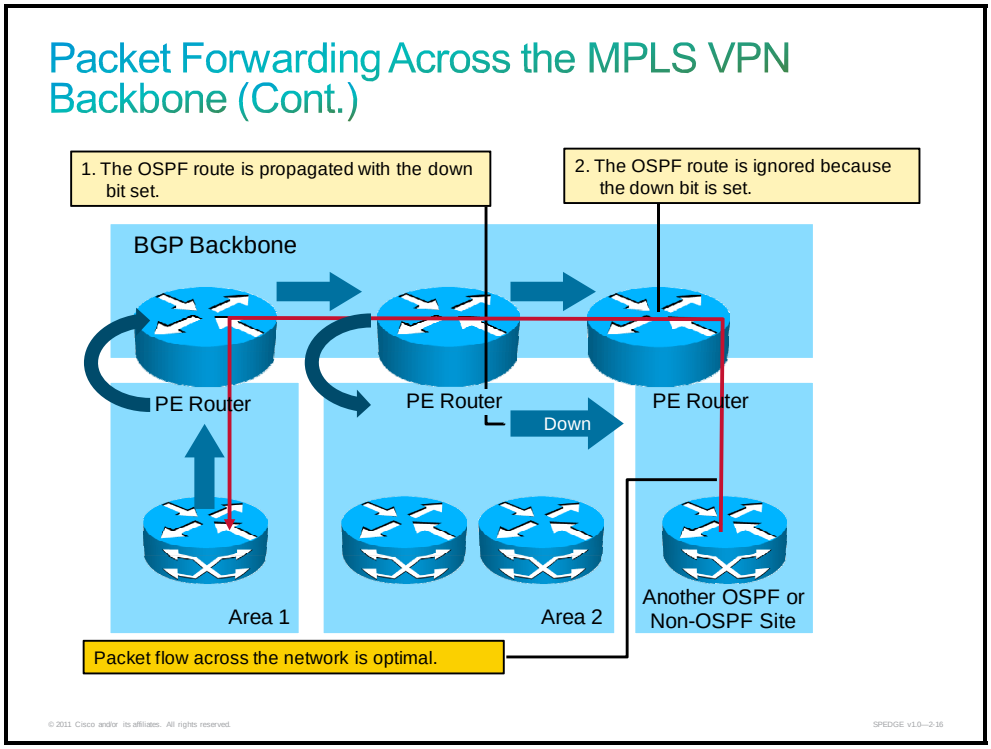
The down bit is used between the PE routers to indicate which routes were inserted into the OSPF topology database from the MPLS VPN superbackbone. Those routes thus will not be redistributed back into the MPLS VPN superbackbone. The PE router that redistributes the MP-BGP route as an OSPF route into the OSPF topology database sets the down bit. The other PE routers use the down bit to prevent this route from being redistributed back into MP-BGP.

https://t.me/learningnets



The OSPF superbackbone implementation with MP-BGP has other implications beyond the potential for routing loops between OSPF and BGP. For example, consider the network in the figure, which shows a typical flow for routing updates.

https://t.me/learningnets



To prevent customer sites from acting as transit parts of the MPLS VPN network, the OSPF route-selection rules in PE routers need to be changed. The PE routers must ignore all OSPF routes with the down bit set, because these routes originated in the MP-BGP backbone, and the MP-BGP route should be used as the optimum route toward the destination.

This rule is implemented with the routing bit in the OSPF LSA. For routes with the down bit set, the routing bit is cleared and these routes never enter the IP routing table, even if they are selected as the best routes by the Shortest Path First (SPF) algorithm.

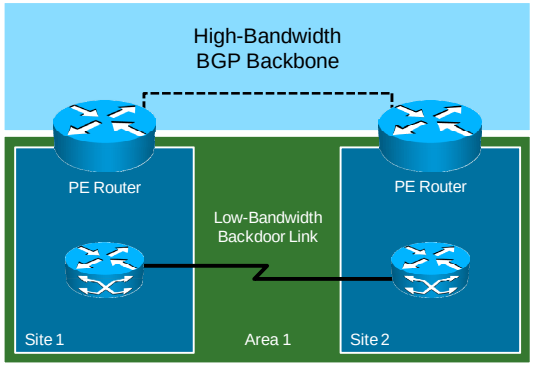
Note	The routing bit is the Cisco extension to OSPF and is used only internally in the router. The routing bit is never propagated between routers in LSA updates.

With the new OSPF route-selection rules in place, packet forwarding in the network that is shown in the figure follows the desired path.

https://t.me/learningnets

Sham Link

- OSPF prefers intra-area paths to interarea paths.
- The path over a backdoor link will always be selected.
- A sham link is a logical intra-area link.
- It is carried by the superbackbone.
- A sham link is required only between two VPN sites that belong to the same area and have a backdoor link for backup purposes.
- OSPF adjacency is established across the sham link.
- LSA flooding occurs across the sham link.



© 2011 Cisco and/or its affiliates. All rights reserved.

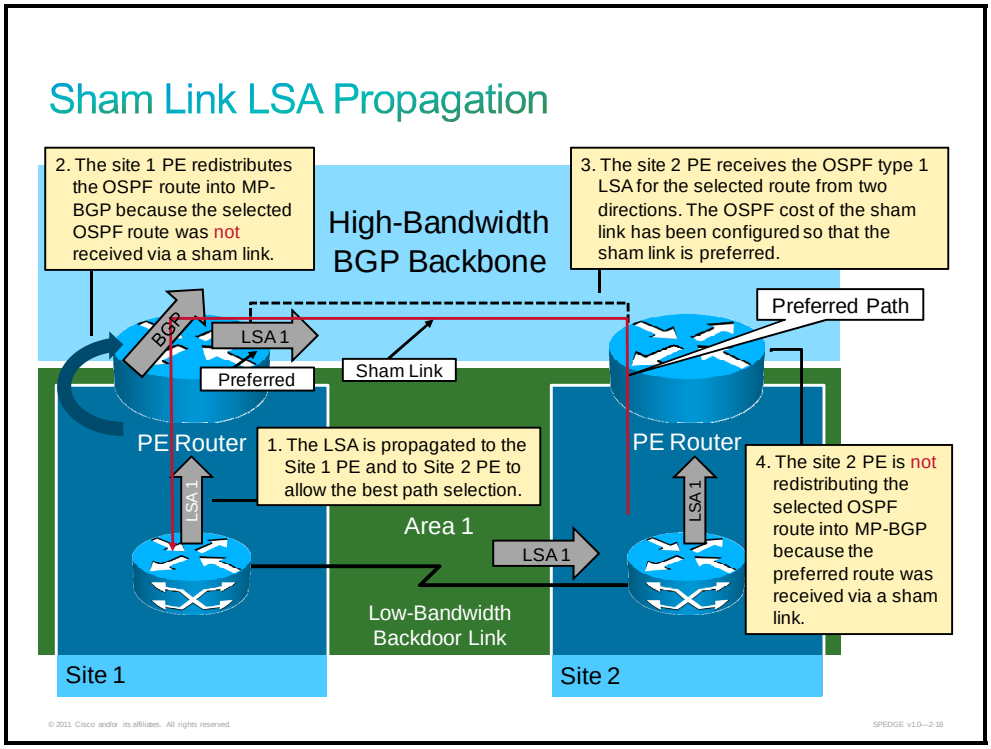
Although OSPF PE-CE connections assume that the only path between two client sites is across the MPLS VPN backbone, backdoor paths between VPN sites may exist.

The figure shows backdoor paths between VPN sites. If these sites belong to the same OSPF area, the path over a backdoor link will always be selected, because OSPF prefers intra-area paths to interarea paths. (PE routers advertise OSPF routes learned over the VPN backbone as interarea paths.) For this reason, OSPF backdoor links between VPN sites must be taken into account so that routing is performed based on policy.

Because each site runs OSPF within the same Area 1 configuration, all routing between the sites follows the intra-area path across the backdoor links rather than over the MPLS VPN backbone.

A sham link is required between any two VPN sites that belong to the same OSPF area *and* share an OSPF backdoor link. If no backdoor link exists between the sites, no sham link is required.

https://t.me/learningnets



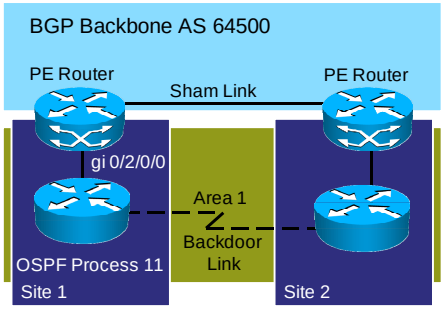
Because the sham link is seen as an intra-area link between PE routers, an OSPF adjacency is created, and a database exchange (for the particular OSPF process) occurs across the link.

After the desired intra-area connectivity is created, the PE routers prefer routes that are learned via the high-bandwidth backbone, because the OSPF cost of the sham link has been configured so that it is preferred over the backdoor link. The implementation results in optimal packet flow.

https://t.me/learningnets

Configuring Sham Links

- A separate /32 address space is required in each PE router for each sham link.
- This /32 address space:
  - Is required so that OSPF packets can be sent over the VPN backbone to the remote end of the sham link
  - Must belong to the VRF
  - Must not be advertised by OSPF
  - Must be advertised by BGP



Cisco IOS and IOS XE

```
Router(config-router)#
area area-id sham-link source-address destination-address cost number
```

Cisco IOS XR

```
RP/0/RP0/CPU0:router(config-ospf)#
vrf vrf-name
 area area-id
  sham-link source-address destination-address cost number
```

© 2011 Cisco and/or its affiliates. All rights reserved.

When you are configuring a sham link, a separate /32 address space is required in each PE router.

These criteria apply to this /32 address space:

- Required so that OSPF packets can be sent over the VPN backbone to the remote end of the sham link
- Must belong to the VRF
- Must not be advertised by OSPF
- Must be advertised by BGP

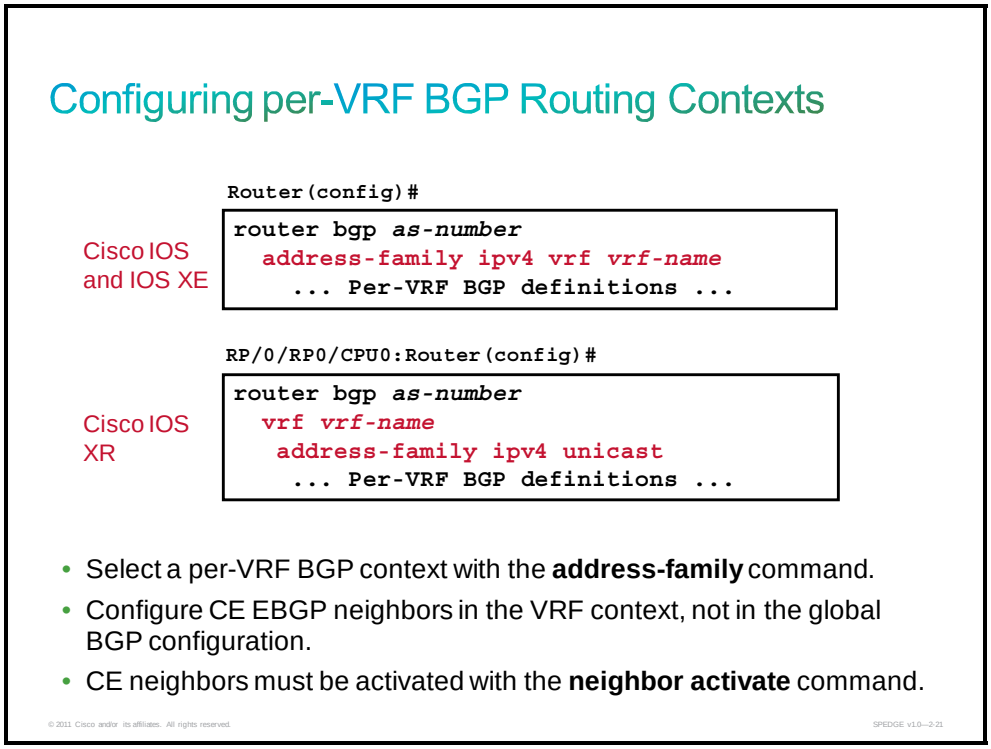
To configure a sham-link interface on a PE router that is running Cisco IOS and IOS XE Software in an MPLS VPN backbone, use the **area sham-link cost** command in global configuration mode.

To configure a sham-link interface using Cisco IOS XR Software, you must first enter VRF configuration mode for OSPF routing and configure an area for the OSPF process using the **area** command. Then you must assign interfaces to the sham link using the **sham-link** command.

https://t.me/learningnets

BGP as the PE-CE Routing Protocol

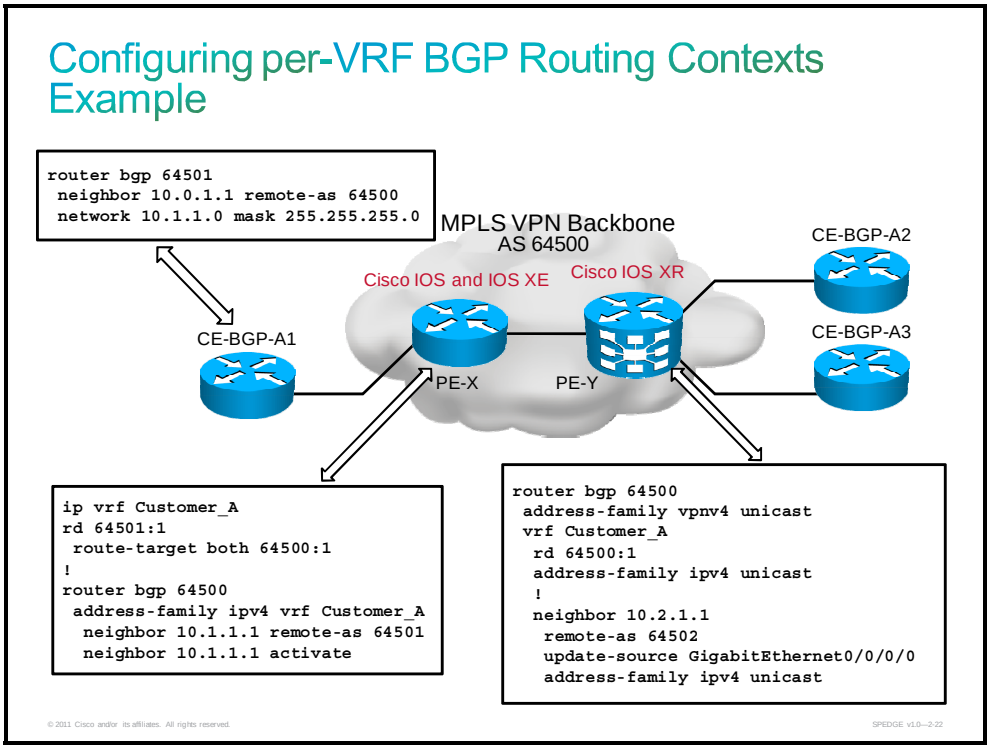
This topic describes BGP as the routing protocol between PE and CE routers.



When you configure BGP as the PE-CE routing protocol, you must start with the per-VRF BGP configuration. Use the **address-family ipv4 vrf** *vrf-name* command in router configuration mode on Cisco IOS and IOS XE devices. Enter address-family configuration mode, and then define and activate the BGP neighbors. You also need to configure redistribution from all other per-VRF routing protocols into BGP.

On Cisco IOS XR devices, first define the VRF with the **vrf** *vrf-name* command in the BGP routing processes. Then select the routing context with the **address-family ipv4 unicast** command. All per-VRF routing protocol parameters (network numbers, passive interfaces, neighbors, filters, and so on) are configured under this address family.

https://t.me/learningnets



The figure shows that BGP is activated on the CE-BGP-A1 router and that the PE-X router is defined as a BGP neighbor. In addition, on the PE-X router, the CE-BGP-A1 router is defined as a BGP neighbor and is activated under the **address-family ipv4 vrf Customer_A** command. Both routers are running Cisco IOS and IOS XE Software.

You can also see that the PE-Y router is running Cisco IOS XR Software. Because all CE routers in the example are running Cisco IOS and IOS XE Software, the configuration for the CE-BGP-A2 and CE-BGP-A3 routers is omitted.

https://t.me/learningnets

Limiting the Number of Routes in a VRF

- Service providers offering MPLS VPN services are at risk of denial-of-service attacks similar to those aimed at service providers offering BGP connectivity:
  - Any customer can generate any number of routes, using resources in the PE routers.
  - Therefore, the resources that are used by a single customer have to be limited.
- Cisco IOS Software offers two solutions:
  - You can limit the number of routes received from a BGP neighbor.
  - You can limit the total number of routes in a VRF.

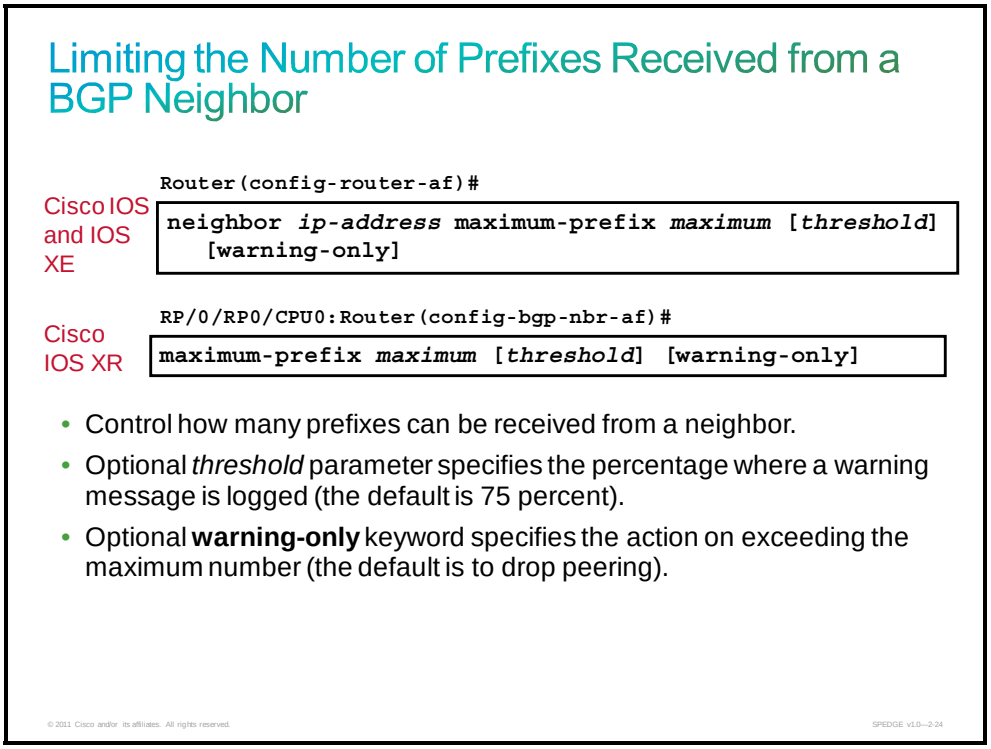
© 2011 Cisco and/or its affiliates. All rights reserved.

MPLS VPN architecture achieves a tight coupling between the customer and the service provider network, resulting in a number of advantages. The tight coupling could also result in a few disadvantages, because the service provider network is exposed to design and configuration errors in customer networks, and a number of new denial-of-service (DoS) attacks are based on routing protocol behavior.

To limit the effect of configuration errors and malicious users, Cisco IOS Software offers two features that limit the number of routes and the resource consumption that are available to a VPN user at a PE router:

- The BGP Maximum-Prefix feature limits the number of routes that an individual BGP peer can send.
- The VRF route limit restricts the total number of routes in a VRF regardless of whether those routes are received from CE routers or from other PE routers via Multiprotocol Internal Border Gateway Protocol (MP-IBGP).

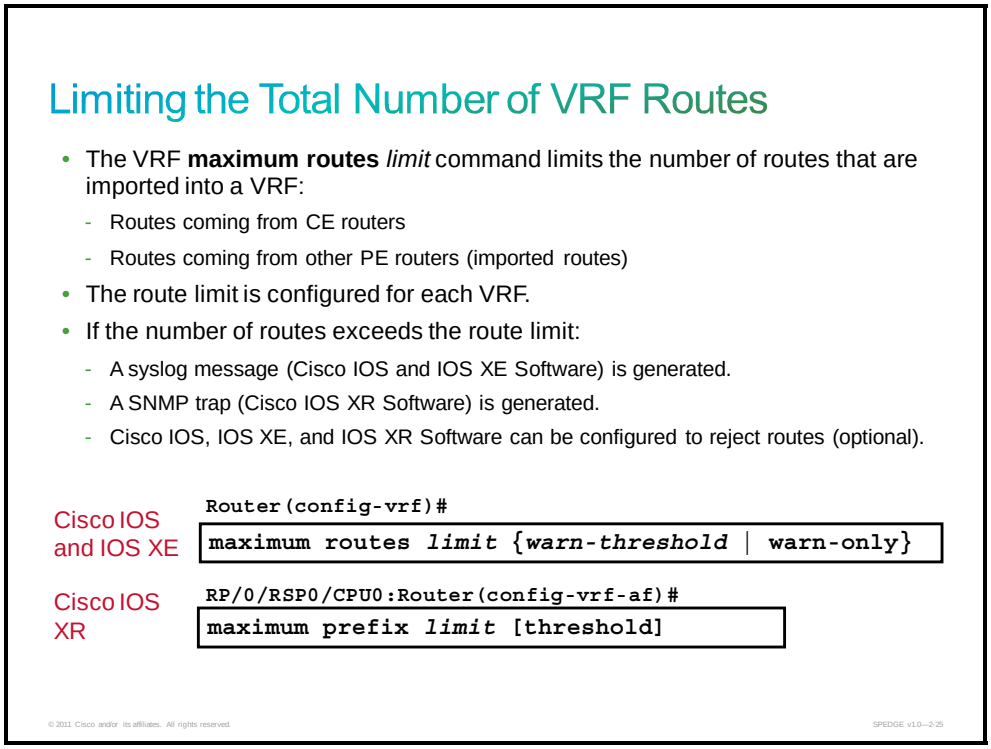
https://t.me/learningnets



To control how many prefixes can be received from a neighbor, use the **maximum-prefix** command for the peer for the appropriate address family.

- Cisco IOS and IOS XE Software:
  neighbor {*ip-address* | *peer-group-name*} **maximum-prefix maximum** [*threshold*] [**restart** *restart-interval*] [**warning-only**]
- Cisco IOS and IOS XE Software:
  maximum-prefix maximum [*threshold*] [**warning-only**]

https://t.me/learningnets



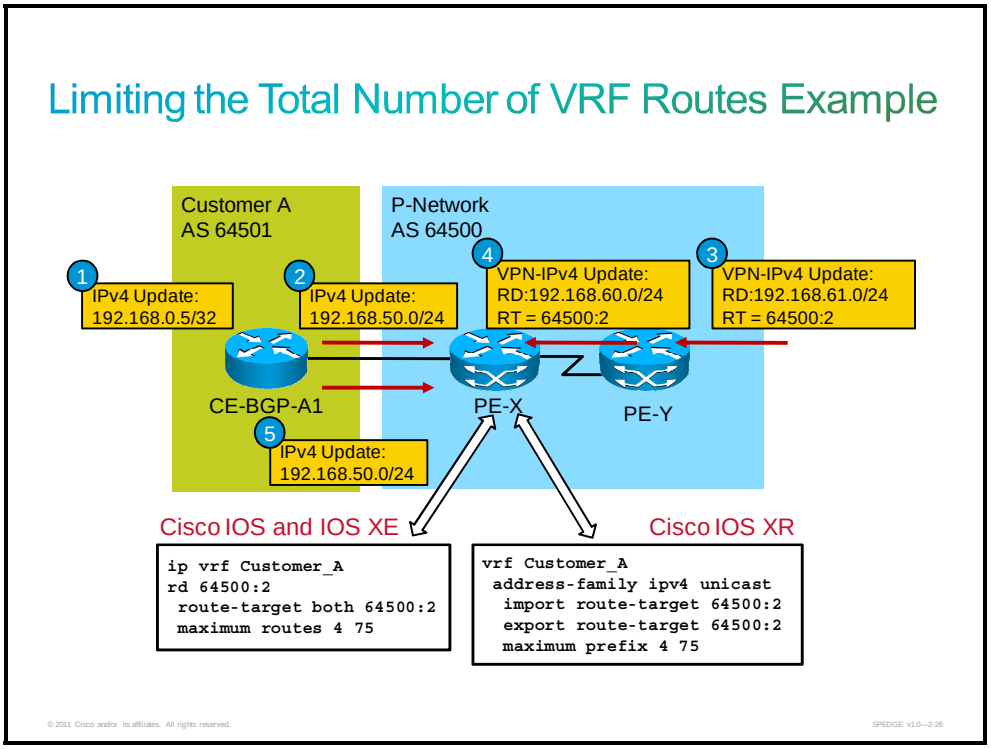
The VRF route limit, unlike the BGP maximum-prefix limit, limits the overall number of routes in a VRF regardless of their origin. As with the BGP Maximum-Prefix feature, the network operator might be warned by a syslog message or Simple Network Management Protocol (SNMP) trap when the number of routes exceeds a certain threshold. Additionally, you can configure Cisco IOS and IOS XE and Cisco IOS XR Software to ignore new VRF routes when the total number of routes exceeds the maximum configured limit.

The route limit is configured for each individual VRF, providing maximum design and configuration flexibility.

Note The per-VRF limit could be used to implement add-on MPLS VPN services. A user wanting a higher level of service might be willing to pay to be able to insert more VPN routes into the network.

To limit the maximum number of routes in a VRF instance to prevent a PE router from importing too many routes, use the **maximum routes** command in VRF configuration mode (IOS and IOS XE Software) or VRF address family configuration mode (IOS XR Software).

https://t.me/learningnets



The network designer can decide to limit the number of routes in a VRF.

In the figure, the network designer has decided to limit the number of routes in a VRF to four, with the warning threshold being set at 75 percent (or three routes).

When the first two routes are received and inserted into the VRF, the router accepts them. When the third route is received, a warning message is generated, and the message is repeated with the insertion of the fourth route.

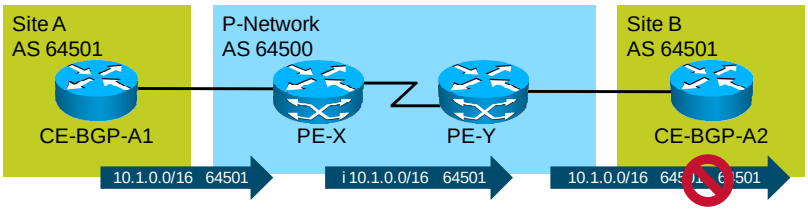
When the PE router receives the fifth route, the maximum route limit is exceeded, and the route is ignored. The network operator is notified through another syslog message.

https://t.me/learningnets

"Same AS Number" Issue

The customer wants to reuse an AS number on several sites:

- CE-BGP-A1 announces network 10.1.0.0/16 to PE-Site-X.
- The prefix announced by CE-BGP-A1 is propagated to PE-Site-Y as an internal route through MP-BGP.
- PE-Site-Y prepends AS 64500 to the AS path and propagates the prefix to CE-BGP-A2.
- CE-BGP-A2 drops the update because AS 64501 is already in the AS path.



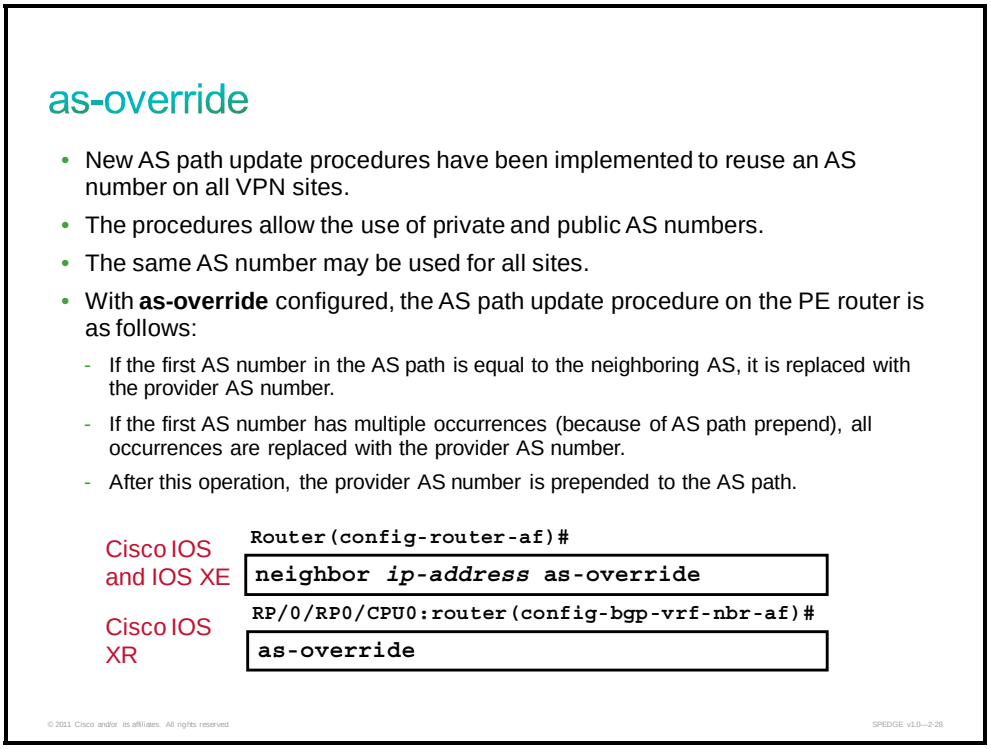
© 2011 Cisco and/or its affiliates. All rights reserved.

Here are the two ways that an MPLS VPN customer can deploy BGP as the routing protocol between PE and CE routers:

- If the customer has previously used any other routing protocol in the traditional overlay VPN network, there are no limitations on the numbering of the customer autonomous systems. Every site can be a separate autonomous system (AS).
- If the customer has used BGP as the routing protocol before, there is a good chance that all the sites (or a subset of the sites) are using the same AS number.

BGP loop-prevention rules disallow discontiguous autonomous systems. Two customer sites with the identical AS number cannot be linked by another AS. If such a setup happens (as in the example in the figure), the routing updates from one site are dropped when the other site receives them. There is no connectivity between the sites.

https://t.me/learningnets



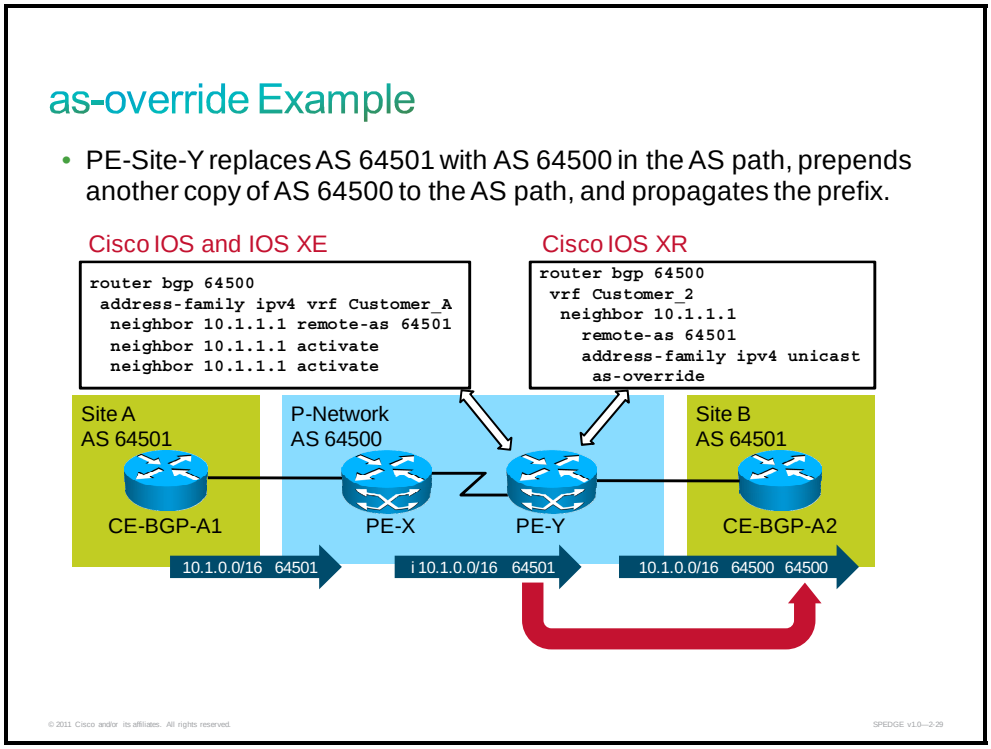
When you are migrating customers from traditional overlay VPNs to MPLS VPNs, it is not uncommon to encounter a customer topology that requires a customer AS number to be used at more than one site. This requirement can cause problems with the loop-prevention rules of BGP. However, the AS-path update procedure in BGP has been modified to address this issue. The new AS-path update procedure supports the use of one AS number at many sites (even between several overlapping VPNs) and does not rely on a distinction between private and public AS numbers.

The modified AS-path update procedure is called *AS override*:

- The procedure is used only if the first AS number in the AS path is equal to the AS number of the receiving BGP router.
- In this case, all leading occurrences of the AS number of the receiving BGP router are replaced with the AS number of the sending BGP router. Occurrences that are farther down the AS path of the AS number of the receiving router are not replaced because they indicate a real routing information loop.
- An extra copy of the sending router AS number is prepended to the AS path. The standard AS number prepending procedure occurs on every External Border Gateway Protocol (EBGP) update.

To configure a PE router to override a site AS number with a provider AS number, use the **as-override** command in the appropriate address family.

https://t.me/learningnets



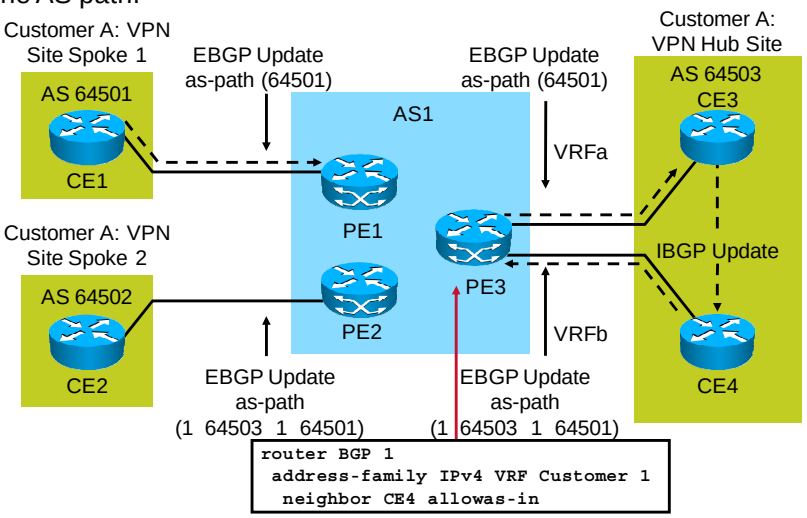
In this figure, customer sites A and B use BGP to communicate with the MPLS VPN backbone. Both sites use AS 64501. Site B would drop the update that was sent by site A without the AS-override mechanism.

The AS-override mechanism, configured on the PE-Site-Y router, replaces the customer AS number (64501) with the provider AS number (64500) before sending the update to the customer site. An extra copy of the provider AS number is prepended to the AS path during the standard EBGP update process.

https://t.me/learningnets

allowas-in

The BGP route is rejected because the PE3 router sees its own AS number in the AS path.

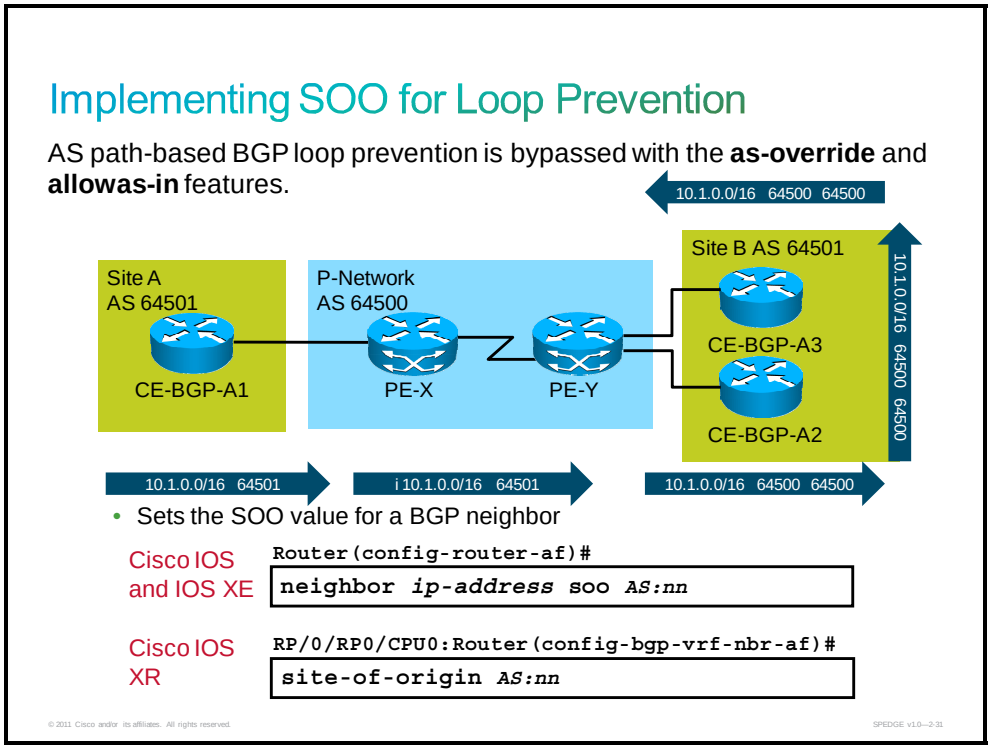


Consider a hub-and-spoke scenario that requires you to permit the routes that are coming from the VRF hub site to re-enter the AS of the service provider. To do so requires that the spoke-to-spoke communication happen through the VRF hub site.

The hub site connects to the provider with two links, which belong to two different VRFs on PE3. One link is used to send updates to the hub site, and one link is used to receive updates from the hub site. For BGP, this setup implies that a route traverses the service provider AS from a VRF spoke site to the VRF hub site and traverses it again on the way to another VRF spoke site. The PE3 router that connects to the VRF hub site sees its own AS number in the AS path, so the BGP route is rejected.

To disable the AS-path loop check, you can configure the command **neighbor allowas-in number** on the PE3 router that connects to the VRF hub site. The **allowas-in** command permits multiple occurrences of the same AS number (in this case, the AS number of the service provider) as the AS number of the BGP speaker in the AS path without BGP denying the route. You can configure a number from 1 to 10 to specify the number of times that the AS number is allowed in the AS path.

https://t.me/learningnets



Most aspects of BGP loop prevention are bypassed when you use either the **as-override** or **allowas-in** command. Routing information loops can still be detected by manually counting occurrences of an AS number in the AS path in an end-to-end BGP routing scenario and then ensuring that the number field in the **allowas-in** command is set low enough to prevent loops.

The ability to continue to detect loops can present a particular problem when BGP is mixed with other PE-CE routing protocols. The Site of Origin (SOO) extended BGP community can be used as an additional loop-prevention mechanism in these situations.

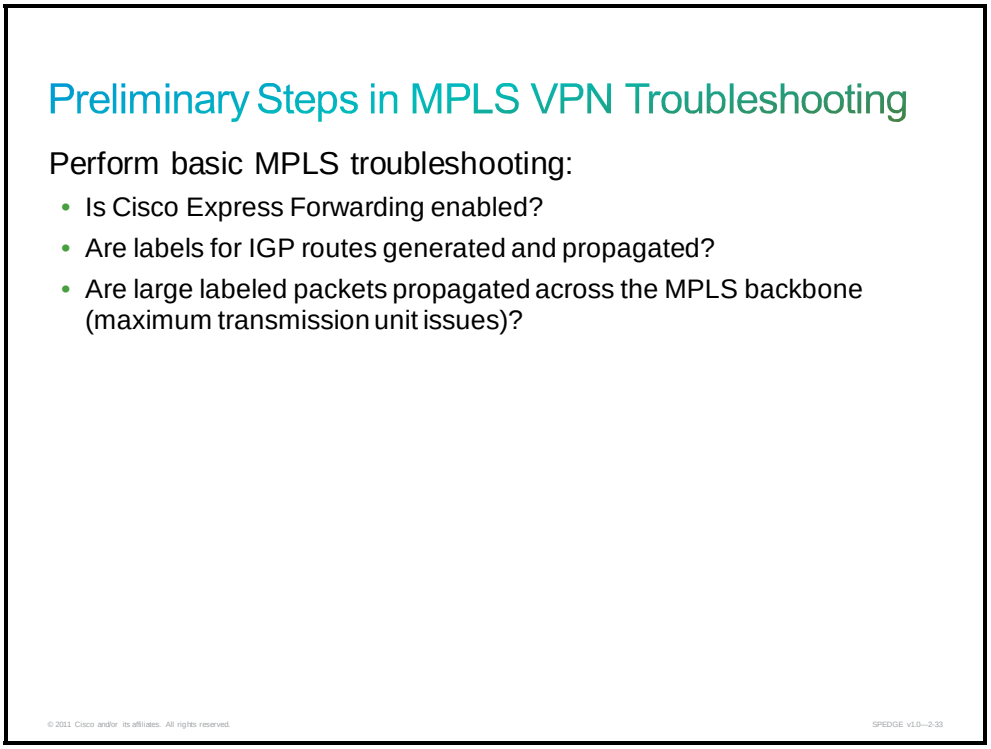
The SOO uniquely identifies the site that originates a route. It is a BGP extended community that prevents routing loops or suboptimal routing, specifically when a back door is present between VPN sites. The SOO provides loop prevention in networks with dual-homed or multihomed sites (sites that are connected to two or more PE routers). You can use it when an IGP is the PE-CE routing protocol. You can also use it when BGP is used between PE and CE, when the AS-path loop prevention cannot be trusted anymore. This situation happens when BGP uses **as-override** or **allowas-in**. If the SOO is configured for a CE router and a VPNv4 route is learned with the same SOO, the route must not be put in the VRF routing table on the PE and advertised to the CE.

Use this command to set the SOO value for a BGP neighbor. The SOO value is set under address-family IPv4 VRF configuration mode either directly for a neighbor or for a BGP peer group.

https://t.me/learningnets

Troubleshooting MPLS VPNs

This topic describes steps that are needed for MPLS VPN troubleshooting.



Before you start in-depth MPLS VPN troubleshooting, you should ask the following standard MPLS troubleshooting questions:

- Is Cisco Express Forwarding enabled on all routers in the transit path between the PE routers?
- Are labels for BGP next hops generated and propagated?
- Are there any maximum transmission unit (MTU) issues in the transit path (for example, LAN switches not supporting a jumbo Ethernet frame)?

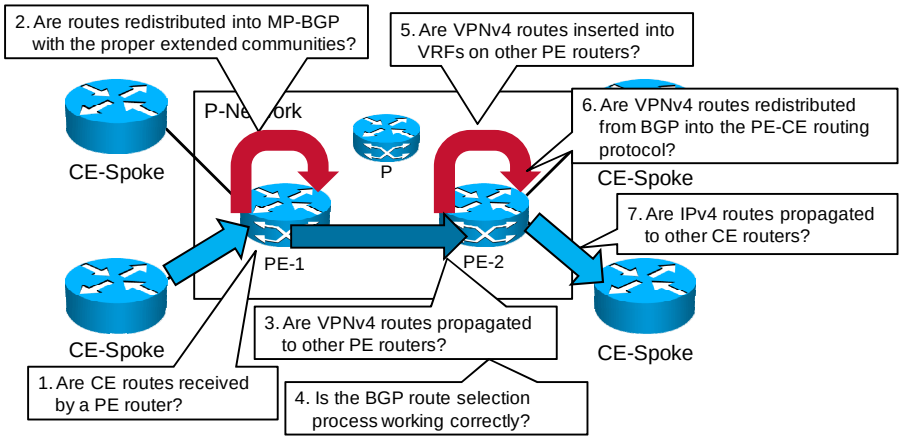
MPLS VPN troubleshooting consists of these two major steps:

- Verifying the routing information flow, using the checks that are outlined in the figure
- Verifying the data flow, or packet forwarding

 © 2012 Cisco Systems, Inc.

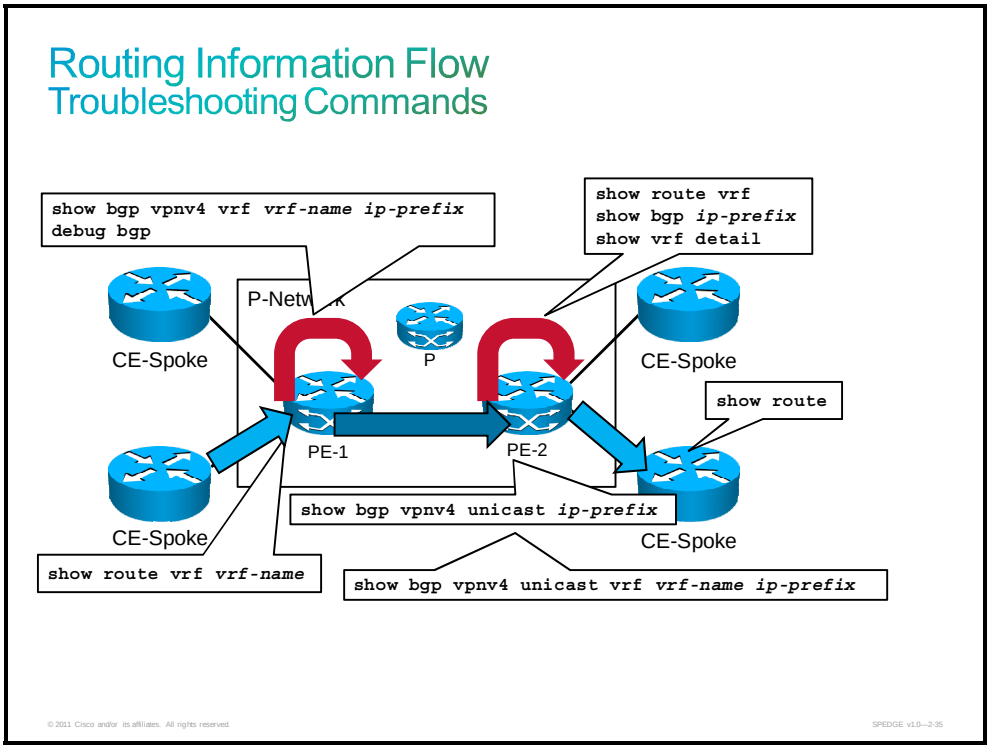
https://t.me/learningnets

Verifying the Routing Information Flow



Verification of the routing information flow should be done systematically, starting at the ingress CE router and moving to the egress CE router.

https://t.me/learningnets



Troubleshooting routing information flow requires the verification of end-to-end routing information propagation between CE routers.

Step 1 The first step is to check the routing information exchange from the CE routers to the PE routers. Use the **show ip route vrf** *vrf-name* (Cisco IOS and IOS XE) or **show route vrf** *vrf-name* (Cisco IOS XR) command to verify that the PE router receives customer routes from the CE router. Use traditional routing protocol troubleshooting if needed.

Step 2 The CE routes that are received by the PE router need to be redistributed into MP-BGP; otherwise, the routes will not be propagated to other PE routers. Proper redistribution of CE routes into a per-VRF instance of BGP can be verified with the **show ip bgp vpnv4 vrf** *vrf-name* (IOS and IOS XE) or **show bgp vpnv4 vrf** *vrf-name* (IOS XR) command. The route distinguisher (RD) prepended to the IPv4 prefix and the route targets (RTs) attached to the CE route can be verified with the **show ip bgp vpnv4 vrf** *vrf-name ip-prefix* (IOS and IOS XE) or **show bgp vpnv4 vrf** *vrf-name ip-prefix* (IOS XR) command.

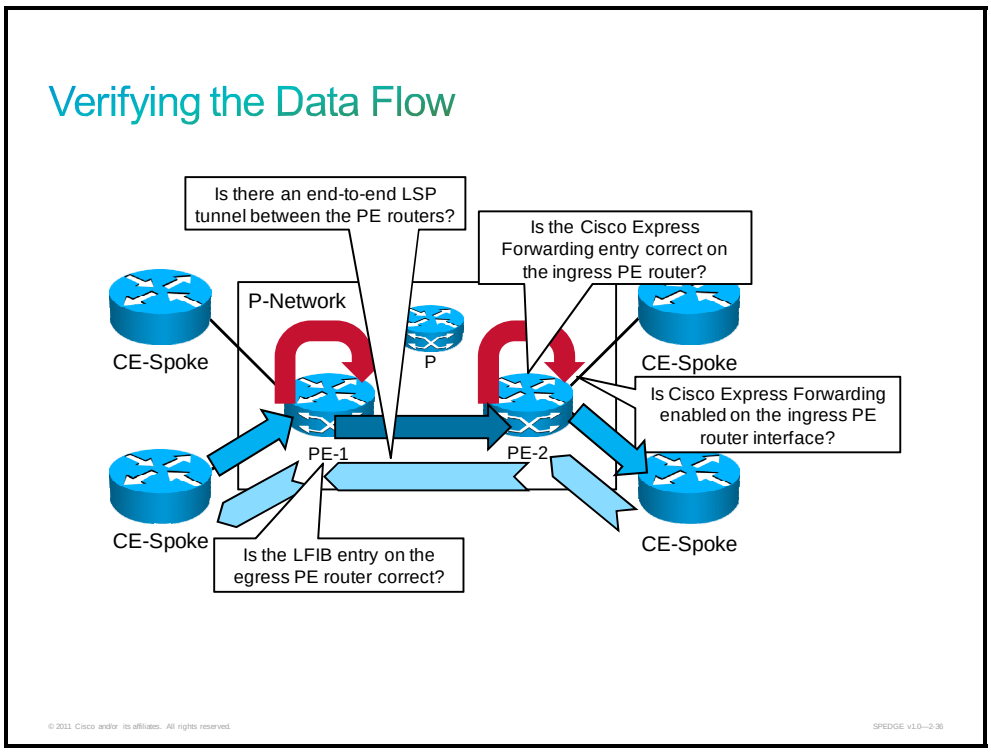
Step 3 The CE routes that are redistributed into MP-BGP need to be propagated to other PE routers. Verify proper route propagation with the **show ip bgp vpnv4 all** *ip-prefix* (IOS and IOS XE) or **show bgp vpnv4 unicast** *ip-prefix* (IOS XR) command on the remote PE router.

Step 4 The VPNv4 routes that are received by the PE router need to be inserted into the proper VRF. This insertion can be verified with the **show ip route vrf** (IOS and IOS XE) or **show route vrf** (IOS XR) command. The validity of the import RTs can be verified with the **show ip bgp vpnv4 all** *ip-prefix* (IOS and IOS XE) or **show bgp vpnv4 unicast** *ip-prefix* (IOS XR) command, which displays the RTs attached to a VPNv4 route. You can also verify the validity of the import RTs with the **show ip vrf detail** (IOS and IOS XE) or **show vrf detail** (IOS XR) command, which lists the import RTs for a VRF. At least one RT attached to the VPNv4 route needs to match at least one RT in the VRF.

https://t.me/learningnets

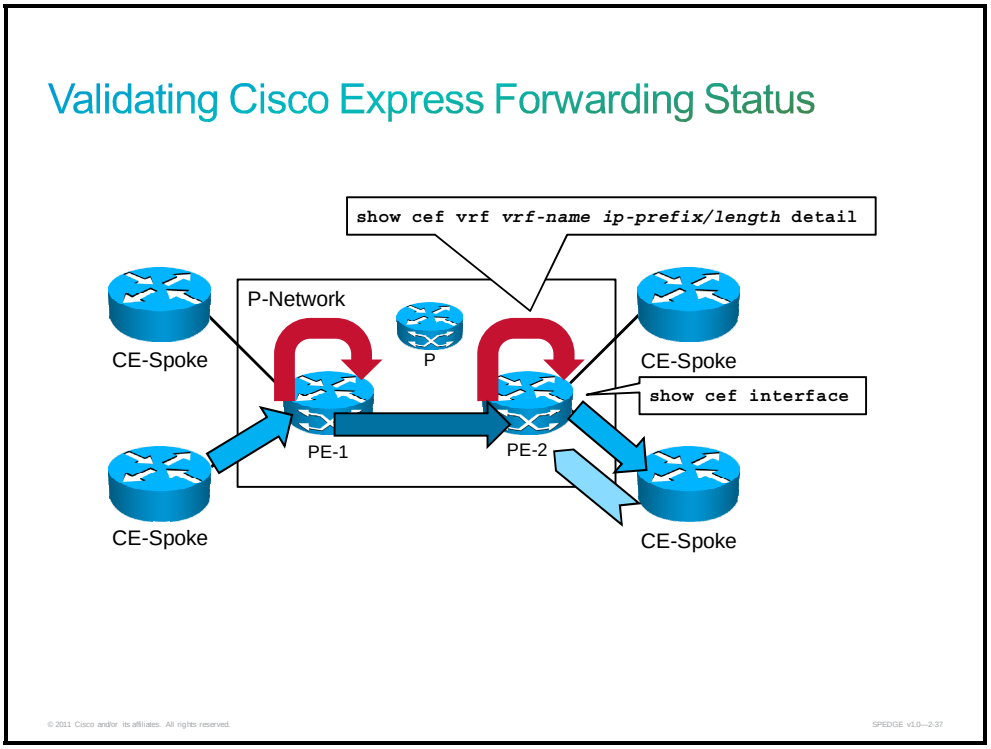
Step 5 Next, the BGP routes that are received via MP-BGP and inserted into the VRF need to be redistributed into the PE-CE routing protocol.

Step 6 Finally, the routes that are redistributed into the PE-CE routing protocol need to be propagated to CE routers.



After you have verified proper route exchange, start MPLS VPN data flow troubleshooting using the checks that are listed in the next figures.

https://t.me/learningnets



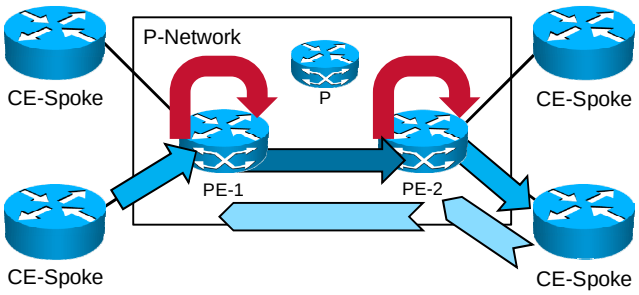
One of the most common configuration mistakes related to data flow is the failure to enable Cisco Express Forwarding on the ingress PE router interface. The presence of Cisco Express Forwarding can be verified with the **show cef interface** command on Cisco IOS, IOS XE, and IOS XR devices.

If Cisco Express Forwarding switching is enabled on the ingress interface, you can verify the validity of the Cisco Express Forwarding entry and the associated label stack with the **show ip cef vrf** *vrf-name ip-prefix* **detail** (IOS and IOS XE) or **show cef vrf** *vrf-name ip-prefix* **detail** (IOS XR) command. The top label in the stack should correspond to the BGP next-hop label as displayed by the **show mpls forwarding-table** (IOS and IOS XE) or **show mpls forwarding** (IOS XR) command on the ingress router. The second label in the stack should correspond to the label allocated by the egress router. You can verify this label by using the **show mpls forwarding-table** (IOS and IOS XE) or **show mpls forwarding vrf** (IOS XR) command on the egress router.

https://t.me/learningnets

Validating the End-to-End Label-Switched Path

- Check for summarization issues. The BGP next hop should be reachable as a host route.
- Quick check—If TTL propagation is disabled, the trace from PE-2 to PE-1 should contain only one hop.
- If needed, check LFIB values hop by hop.
- Check for MTU issues on the path. MPLS VPN requires a larger label header than pure MPLS.

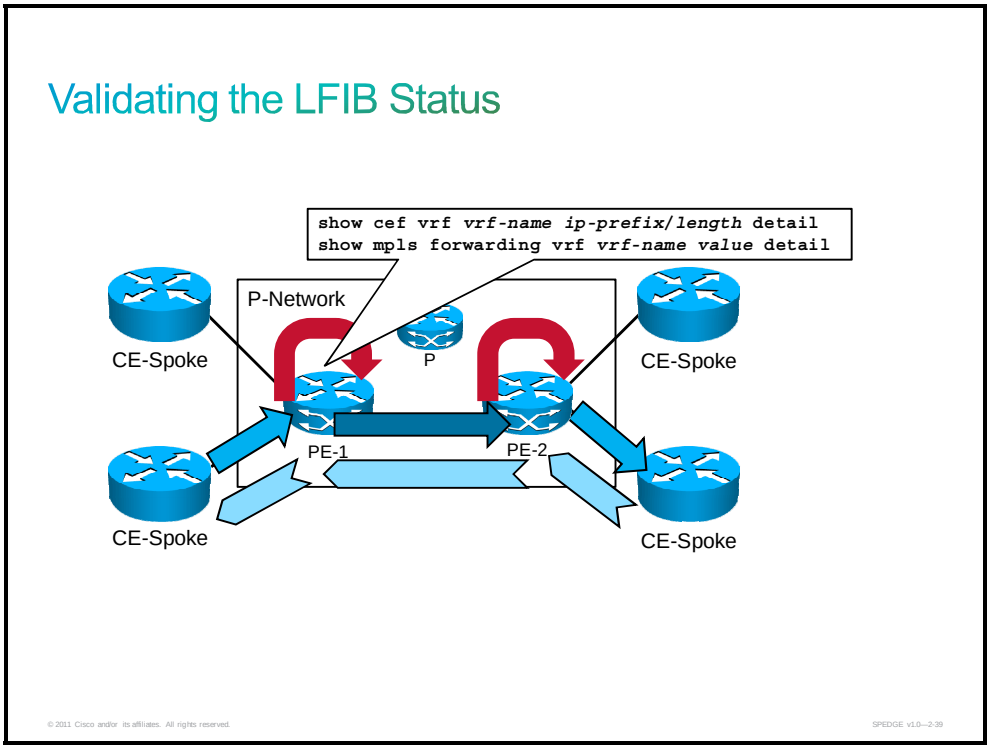


© 2011 Cisco and/or its affiliates. All rights reserved.

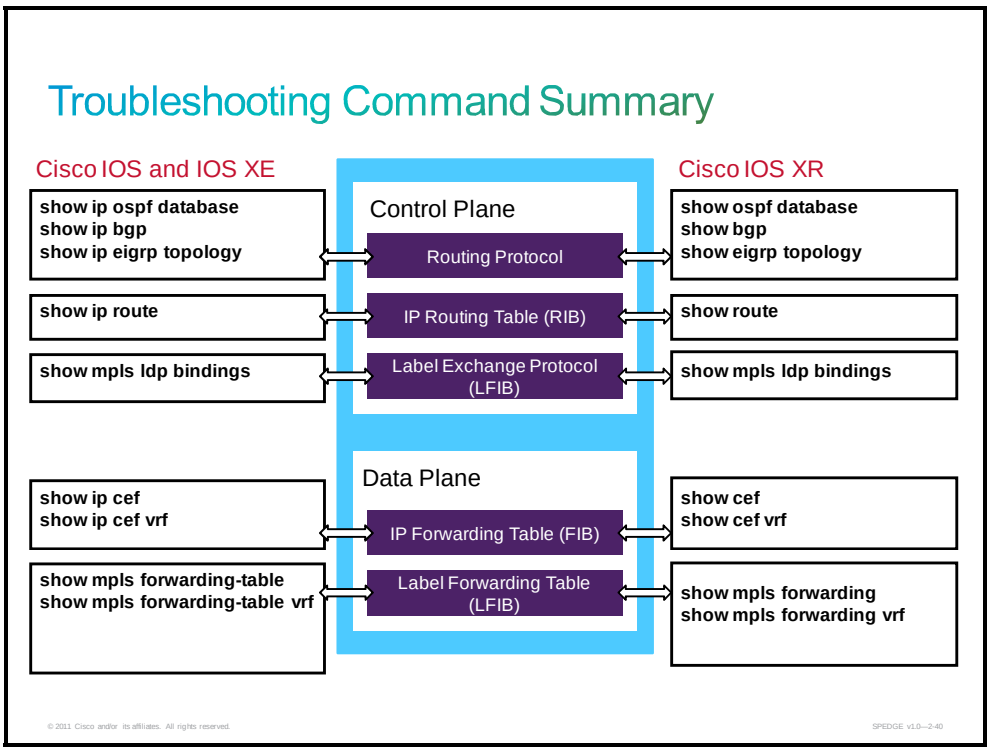
If Cisco Express Forwarding is enabled on the ingress interface and the Cisco Express Forwarding entry contains the proper labels, the data flow problem might lie inside the MPLS core. Two common mistakes include summarization of BGP next hops inside the core IGP and MTU issues.

The quickest way to diagnose summarization problems is to disable IP Time to Live (TTL) propagation into the MPLS label header using the **no mpls ip ttl-propagate** (IOS and IOS XE) or **mpls ip ttl-propagate disable** (IOS XR) configuration command on the provider router (P router) and PE routers. The **traceroute** command from the ingress PE router toward the BGP next hop should display no intermediate hops when TTL propagation is disabled. If intermediate hops are displayed, the label-switched path (LSP) tunnel between PE routers is broken at those hops, and the VPN traffic cannot flow.

https://t.me/learningnets



As a last troubleshooting measure (usually not needed), you can verify the contents of the label forwarding information base (LFIB) on the egress PE router and compare them with the second label in the label stack on the ingress PE router. A mismatch indicates an internal Cisco IOS Software error that you will need to report to the Cisco Technical Assistance Center (TAC).



The figure shows the relevant commands for Cisco IOS, IOS XE, and Cisco IOS XR devices to troubleshoot the control plane and the data plane of an MPLS router.

https://t.me/learningnets

Summary

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

Summary

- OSPF as a PE-CE routing protocol is implemented as a separate routing process.
- BGP is very scalable and predictable as a PE-CE routing protocol.
- MPLS VPN troubleshooting has two main steps: verifying routing information flow and verifying proper data flow.

© 2011 Cisco and/or its affiliates. All rights reserved.

https://t.me/learningnets

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 4

Deploying IPv6 and MPLS

Overview

When deploying IPv6, you need to understand how to deploy it over a Multiprotocol Label Switching (MPLS) network. This lesson introduces some concepts that can be used when you design or deploy IPv6 in your environment.

Objectives

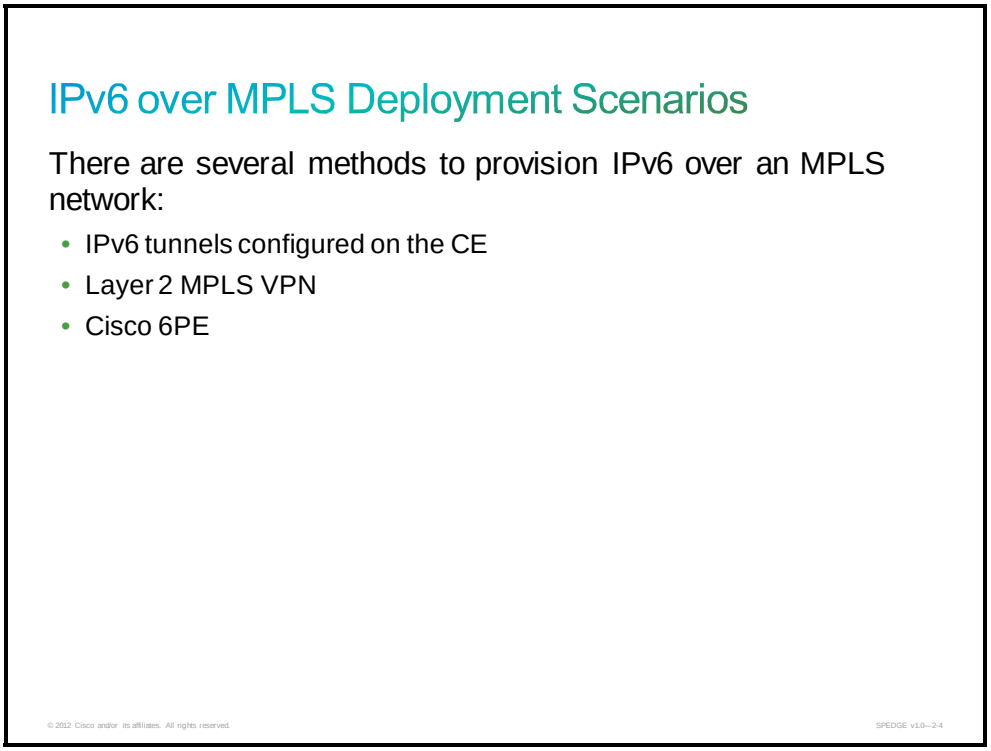
Upon completing this lesson, you will be able to describe how to configure routing protocols between PE and CE routers. You will be able to meet this objective:

- Describe the various methods that are used to deploy IPv6 over MPLS

https://t.me/learningnets

Support for IPv6 in MPLS

This topic describes the various methods that are used to deploy IPv6 over MPLS.

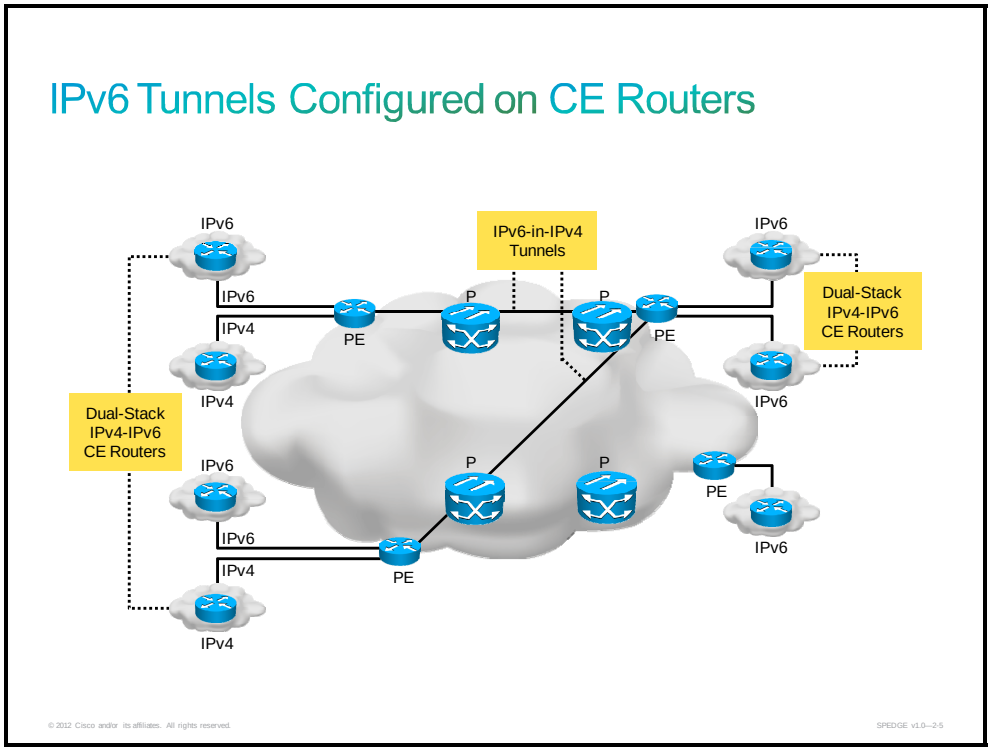


IPv6 over MPLS backbones enables isolated IPv6 domains to communicate with each other over an MPLS IPv4 network. Depending on the deployment scheme that is chosen and on the current MPLS implementation, enabling IPv6 services might require fewer backbone infrastructure upgrades and less reconfiguration of core routers than enabling IPv4 services, because forwarding is based on labels rather than on the IP header, thus providing a very cost-effective strategy for the deployment of IPv6.

Additionally, the inherent VPN and traffic engineering services that are available within an MPLS environment allow IPv6 networks to be combined into VPNs or extranets over an infrastructure that supports IPv4 VPNs and MPLS terminal equipment. A variety of deployment strategies are available or under development:

- **Deploying IPv6 by using tunnels on the customer edge (CE) routers:** This strategy has no impact on and requires no changes to the MPLS provider (P) routers or provider edge (PE) routers. The strategy uses IPv4 tunnels that are terminated on the CE routers to encapsulate the IPv6 traffic, which then appears as IPv4 traffic within the MPLS provider network.
- **Deploying IPv6 over a Layer 2 MPLS VPN:** This strategy, applicable only on specific Cisco routers (such as the Cisco 12000 and 7600 Series Routers), also requires no changes to the core routing mechanisms, assuming that the provider is already capable of supporting Layer 2 MPLS VPNs.
- **Deploying IPv6 on the PE routers via Cisco IPv6 Provider Edge Router over MPLS (6PE):** This strategy requires changes to the PE routers to support a dual-stack implementation, but all the core functions remain IPv4-only, and so the P routers require no changes.

https://t.me/learningnets



Using tunnels on the CE routers is the simplest way to deploy IPv6 over MPLS networks. This method has no impact on the operation or infrastructure of MPLS and requires no changes to the P routers in the core or to the PE routers that connect to customers.

IPv6 tunnels that are configured on a CE have these characteristics:

- The CE routers are IPv6-aware (dual stack).
- A mesh of IPv6-over-IPv4 tunnels connects CE to CE.
- Overhead includes the IPv4 header and MPLS header.
- MPLS and VPN support native IPv4 and IPv6 tunnels.

Communication between the remote IPv6 domains uses standard tunneling mechanisms, running IPv6 over IPv4 tunnels, as the way that MPLS VPNs support native IPv4 tunnels. The CE routers need to be upgraded to dual stack and configured by using manually configured tunnels or 6to4 tunnels. However, communication with the PE routers is IPv4-only, and the traffic appears to be IPv4 to the MPLS domain. The dual-stack routers use the 6to4 tunnel addresses or an IPv6 prefix that is assigned by a virtual IPv6 ISP (possibly other than the MPLS provider). The figure shows an example for deployment of IPv6 by using tunnels on the CE routers.

This alternative is attractive because it has no impact on the MPLS P or PE routers, because it uses IPv4 tunnels to encapsulate IPv6 traffic. After being encapsulated, the traffic appears as IPv4 traffic within the network. However, this option does not permit the delegation of a /48 prefix address from the MPLS service provider address space to customers. Nonetheless, it is probably useful for specific applications, such as adding IPv6 services to an MPLS VPN by configuring the tunnel on the CE. All Cisco IOS Software tunneling mechanisms support this capability.

https://t.me/learningnets

Pros

From the provider standpoint, using tunnels on CE routers is the simplest way to deploy IPv6 over MPLS networks. This approach has no impact on the operation or infrastructure of MPLS and requires no changes to either the P routers in the core or the PE routers that connect to customers.

Cons

Using tunnels on CE routers presents the same problems as any configured tunnel scheme. As the number of tunnel partners grows, the number of tunnels in an any-to-any mesh increases quickly, making management of the tunnels error-prone and expensive. This problem is not specific to MPLS but rather is a general issue with manually configured tunnels. Automatic tunnels (such as 6to4) can be used to reduce the effort that is required to manage tunnels on the CE, but automatic tunnels have other issues that must be considered.

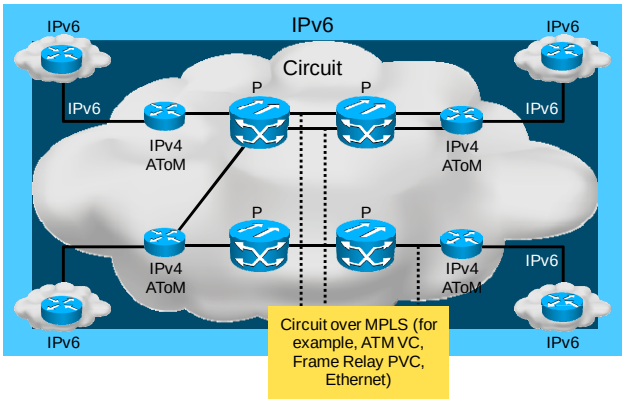
Delegating a global IPv6 prefix for an ISP is also difficult because the ISP does not furnish the IPv6 connectivity; it merely provides transport. Even when the ISP does allocate an address, early traffic levels will be modest, and later traffic might shift from IPv4 to IPv6. Therefore, there is not necessarily a net increase in carried packets. Because the service provider provides transport only over IPv4 over MPLS, this scenario offers little possibility of generating additional revenue for the provider.

From the provider standpoint, using tunnels on CE routers is not a good mechanism to carry IPv6 traffic across an IPv4-based MPLS network. The provider basically puts responsibility for IPv6 transit on the end customers, because the MPLS provider recognizes only IPv4 traffic.

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Layer 2 MPLS VPNs



© 2012 Cisco and/or its affiliates. All rights reserved.

Using any circuit transport to deploy IPv6 over MPLS networks has no impact on the operation or infrastructure of MPLS. Circuit transport requires no changes to either the P routers in the core or the PE routers that connect to customers, assuming that the provider has already implemented circuit-over-MPLS technology. This design is a Layer 2 solution.

IPv6 using a Layer 2 MPLS VPN has these characteristics:

- IPv6 is run over pseudowires (see the IETF document draft-ietf-pwe3-wildcard-pw-type-02.txt).
- Edge MPLS routers need to support a Layer 2 MPLS VPN.
- A mesh of virtual circuits connects PE to PE.
- The PE routers are regular IPv4 routers that are used to aggregate customer IPv6 routers.

Communication between remote IPv6 domains runs native IPv6 protocols over a dedicated link, in which the underlying mechanisms are fully transparent to IPv6. The IPv6 traffic is tunneled by using a Layer 2 MPLS VPN or Ethernet over MPLS (EoMPLS), with the IPv6 routers connecting through an ATM or Ethernet interface, respectively. The PE will require an upgrade to a Layer 2 MPLS VPN if it does not already contain that support. The figure shows an example of IPv6 deployment over any circuit transport over MPLS.

The alternative of IPv6 over Layer 2 MPLS VPN is likely to be used by service providers that have ATM or Ethernet links to CE routers. This alternative is fully transparent to users. However, like the previous option, this option does not permit the delegation of a /48 prefix address from the service provider address space to customers. In addition, IPv6 over Layer 2 MPLS VPN does not scale as well as 6PE because service providers need to create circuits through their MPLS networks for their customers.

https://t.me/learningnets

Pros

There are some advantages to using IPv6 over a Layer 2 MPLS VPN:

- Simple solution for service providers with ATM or Ethernet links to CE routers and circuit over MPLS capability
- Fully transparent IPv6 communication, from the provider perspective
- No required reconfiguration of the MPLS core

Cons

There are also some disadvantages to using IPv6 over a Layer 2 MPLS VPN:

- Equivalent to Layer 2 VPN complexity
- The need to upgrade PE routers that connect to customers to a Layer 2 MPLS VPN
- The difficulty of delegating a global IPv6 prefix for an ISP

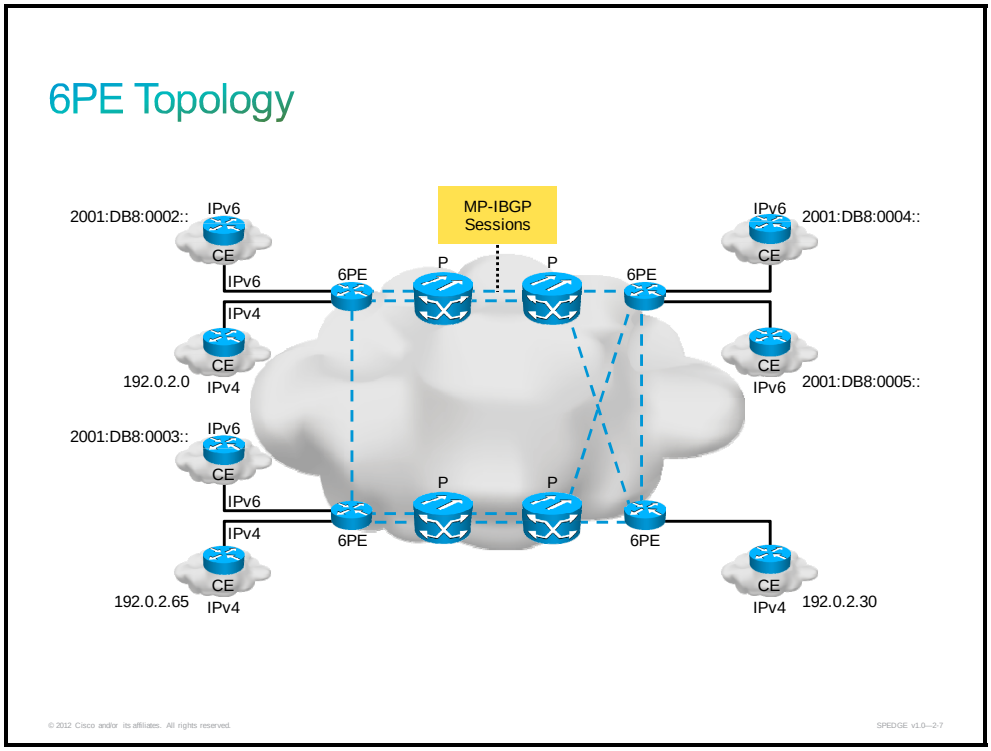
Description of Pros and Cons

Additional points about IPv6 over Layer 2 MPLS VPN are as follows:

- The use of a Layer 2 MPLS VPN to deploy IPv6 over MPLS networks does not affect the operation or infrastructure of MPLS. Neither the P router software version nor the configuration is changed. However, PE routers that connect to customers must be upgraded to Layer 2 MPLS VPN. Delegating a global IPv6 prefix for an ISP is also difficult.
- Communication between remote IPv6 domains runs native IPv6 protocols over a dedicated link in which the underlying mechanisms are fully transparent to IPv6. The IPv6 traffic is tunneled via a Layer 2 MPLS VPN, with the IPv6 routers connected through an ATM or Ethernet interface. This potentially heavy tunnel meshing is a drawback of the solution.

 © 2012 Cisco Systems, Inc.

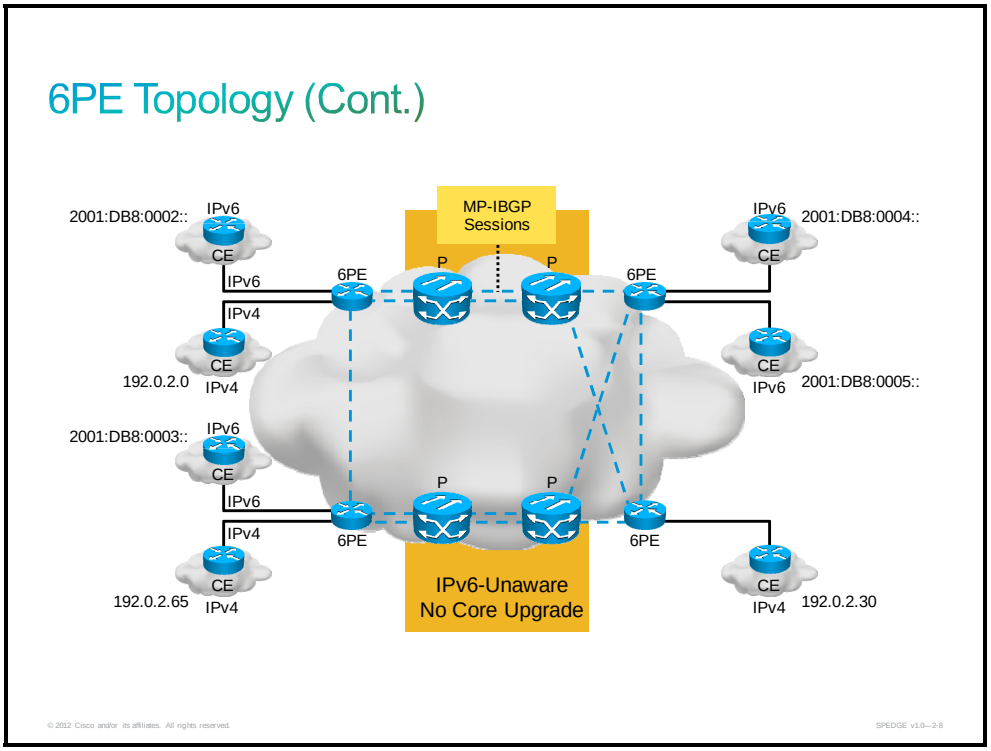
https://t.me/learningnets



Cisco 6PE enables IPv6 islands to communicate with each other over an MPLS IPv4 core network by using MPLS label-switched paths (LSPs). The method relies on Border Gateway Protocol (BGP) extensions in the IPv4 network PE routers (Cisco 6PE routers) to exchange IPv6 reachability information and an MPLS label for each IPv6 address prefix that is announced. Cisco 6PE routers are dual stack (IPv4 and IPv6). This design is a Layer 3 solution.

For the IPv6 transport to be transparent to all but Cisco 6PE routers, you must impose a hierarchy of labels at the Cisco 6PE ingress router. The top label provides connectivity inside the IPv4 MPLS core network. Label Description Protocol (LDP), Tag Distribution Protocol (TDP), or Resource Reservation Protocol (RSVP) distributes this label. The inner label is used for IPv6 forwarding at each Cisco 6PE egress router. This label is distributed by Multiprotocol Border Gateway Protocol (MP-BGP) in the VPN-IPv6 address family.

https://t.me/learningnets



Cisco 6PE uses MPLS to minimize changes in the ISP core network.

With Cisco 6PE, IPv6 forwarding in the core of the MPLS network is done by label switching, eliminating the need for either IPv6-over-IPv4 tunnels or an additional Layer 2 encapsulation. It allows the appearance of a native IPv6 service to be offered across the network.

Each PE router that must support IPv6 connectivity needs to be upgraded to dual stack (becoming a Cisco 6PE router) and configured to run MPLS on the interfaces that connect to the core. Depending on the site requirements, each router can be configured to forward IPv6 or IPv6 and IPv4 traffic on the interfaces to the CE routers, thus providing the ability to offer only native IPv6 or both IPv6 and native IPv4 services. The Cisco 6PE router exchanges either IPv4 or IPv6 routing information with the CE router through any of the supported routing protocols (depending on the connection). It switches IPv4 and IPv6 traffic over the native IPv4 and IPv6 interfaces that do not run MPLS.

The Cisco 6PE router uses MP-BGP to exchange reachability information with the other 6PE routers in the MPLS domain. It shares a common IPv4 routing protocol (such as Open Shortest Path First [OSPF] or Intermediate System-to-Intermediate System [IS-IS]) with the other P and PE devices in the domain. Interior gateway protocol (IGP) is used for deriving next-hop information.

The Cisco 6PE routers encapsulate IPv6 traffic by using two levels of MPLS labels. The top label is distributed by an LDP or TDP that devices in the core use to carry the packet to the egress Cisco 6PE router, using IPv4 routing information. The second or bottom label is associated with the IPv6 prefix of the destination through MP-BGP version 4 (MP-BGP4).

Multiprotocol Extensions for BGP4

Multiprotocol extensions for BGP4 are characterized as follows:

- BGP4 is the primary mechanism to exchange IPv6 reachability information for Cisco 6PE and Cisco IPv6 VPN Provider Edge Router over MPLS (6VPE).
- Label binding information is piggybacked with the prefix advertisement in a Network Layer Reachability Information (NLRI) attribute.

https://t.me/learningnets

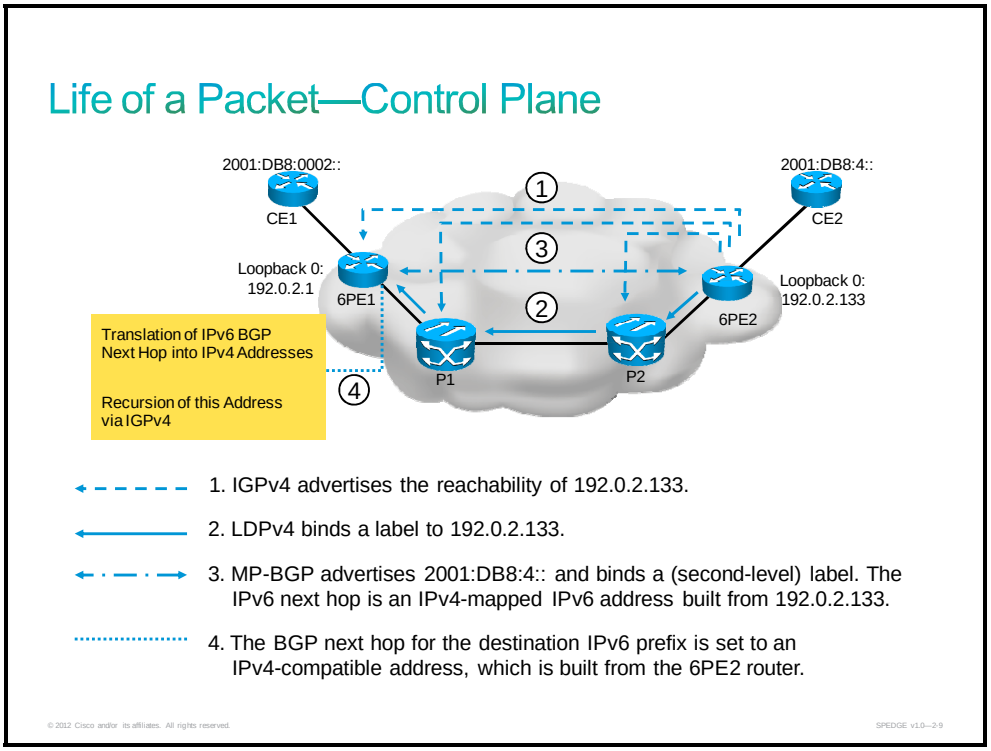
- BGP protocol extensions, which carry routing information about other protocols, include multicast, MPLS, IPv6, VPN-IPv4, labeled IPv6 unicast (Cisco 6PE), and Cisco 6VPE.
- The exchange of multiprotocol NLRI must be negotiated at the session setup. The exchange uses BGP capability advertisement negotiation procedures.

MP-BGP is the supported exterior gateway protocol (EGP) for IPv6. MP-BGP extensions for IPv6 support the same features and functionality as IPv4 BGP. IPv6 enhancements to MP-BGP include support for an IPv6 address family, NLRI, and next-hop attributes that use IPv6 addresses.

The Cisco 6PE multipath feature uses Multiprotocol Internal BGP (MP-IBGP) to distribute IPv6 routes over the MPLS IPv4 core network and to attach an MPLS label to each route.

Note See RFC 2858, *Multiprotocol Extensions for BGP4* for more information. (RFC 2283 is obsolete.)

https://t.me/learningnets



This example describes how IPv6 transport is achieved across an IPv4-based MPLS provider network that is running an IPv4 IGP such as OSPF version 2 (OSPFv2), and how reachability is established between Cisco 6PEs.

1. IGPv4 advertises the reachability of PE routers.

The Cisco 6PE routers exchange IPv6 routing information through MP-BGP. The Cisco 6PE routers peer together through MP-BGP4 to exchange IPv6 reachability with the MPLS cloud and to perform IPv6 forwarding. IGP advertises the reachability of the BGP next-hop address; that is, 192.0.2.133.

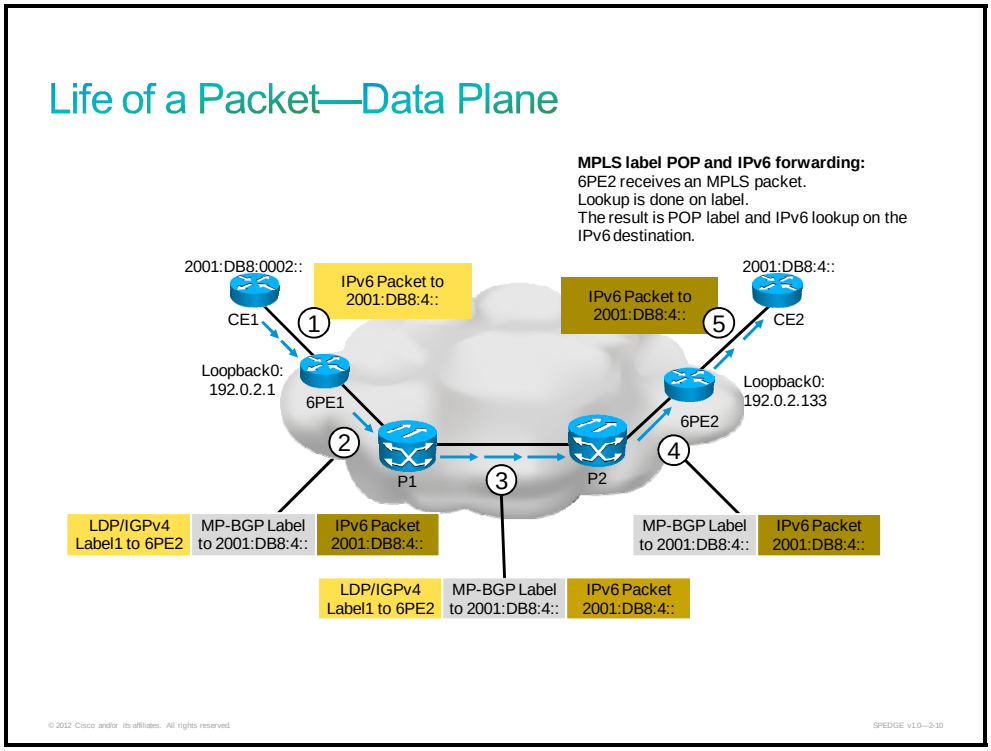
2. LDPv4 binds labels to create LSPs.

3. The 6PE2 router sends MP-BGP advertisements to the ingress PE router, 6PE1.

4. The BGP next hop for the destination IPv6 prefix is set to an IPv4-compatible address, which is built from the 6PE2 router.

The MPLS label binds to the IPv4- or IPv6-mapped address, which is formed from the IPv4 address of 6PE2 (192.0.2.133). Thus, other Cisco 6PE routers (which are part of the mesh edges around the MPLS core) see 6PE2::FFFF:192.0.2.133 as the next-hop IPv6 address when a packet needs to be forwarded to the 2001:DB8:4:: network.

https://t.me/learningnets



This example describes how routers forward packets in data plane:

1. An IPv6 packet arrives at the ingress Cisco 6PE router.

An IPv6 packet that is destined to a network with prefix 2001:DB8:4:: enters through 6PE1. The MPLS core must be used to carry the packet to its destination, via 6PE2. The routing configuration via IBGP has already advertised the next-hop IPv6 address (::FFFF:192.0.2.133) to 6PE1.

2. The ingress Cisco 6PE router binds MP-BGP (inner) and LDP-IGPv4 (top, or outer) labels to the packet.

6PE1 does a lookup on the IPv6 prefix. Two labels are then created and bound to the packet. MP-BGP binds 2001:DB8:4:: and LDP-IGPv4 binds its label to the IPv4 address of 6PE2 (192.0.2.133). In this way, MP-BGP can carry either IPv4 packets or IPv6 packets from one Cisco 6PE router to another.

3. The IPv6 packet is forwarded to the first P router.

MPLS label switching is IPv6-unaware. When P1 receives the MPLS packet, it looks only at the LDP-distributed MPLS label, and a lookup is performed. The result is the imposition of Label2, and the MPLS packet is forwarded to P2. The IPv6 traffic is transparent to the core of the MPLS network.

4. The IPv6 packet is forwarded to the next P router.

P2 receives an MPLS packet, which performs an MPLS lookup on Label2 and results in the point of presence (POP) label. As a result, the packet is forwarded to the IPv4 address to which it is bound. The packet still has the inner MPLS label, which is required because a raw IPv6 packet is not allowed within the MPLS core.

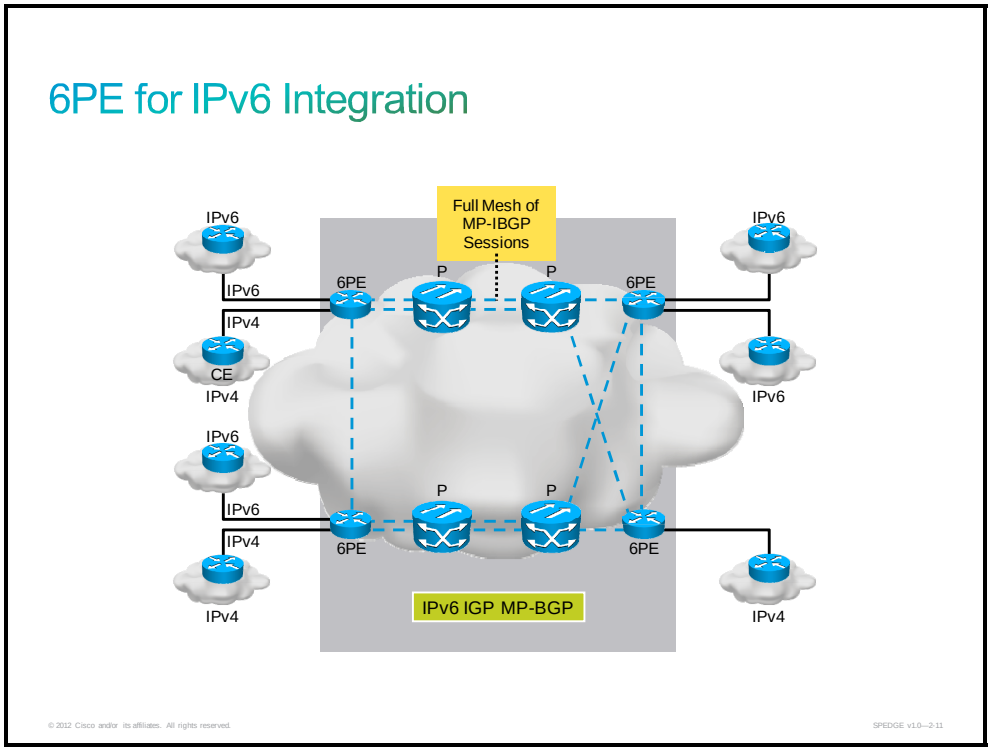
https://t.me/learningnets

5. The egress PE routers remove the MPLS labels and forward the IPv6 packet.

The 6PE2 router examines the MPLS packet. The result is that the remaining MPLS label is removed and an IPv6 lookup on the IPv6 destination address is performed, using normal IP destination-based routing. The Cisco 6PE router forwards the IPv6 packet on to the CE device, based on the results of the IPv6 lookup on the IPv6 destination address.

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets



Using Cisco 6PE for IPv6 integration has these benefits:

- A IPv6 CE has only one PE routing peer, regardless of how many remote IPv6 CEs it communicates with.
- No changes are required on an IPv6 CE when remote CEs are added or removed. Reachability is automatically learned via MP-BGP.
- No tunnel or circuit configuration is required.
- Cisco 6PE offers a scalable and flexible solution. Benefits are analogous to the Layer 3 VPN solution for IPv4, in RFC 2547bis.

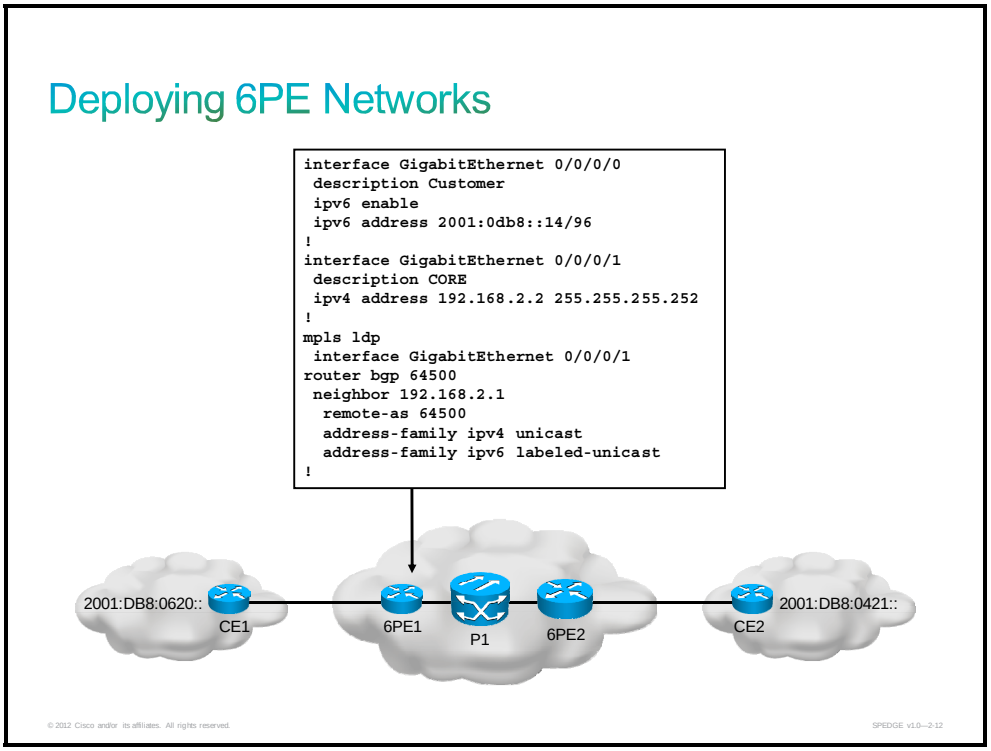
Service providers that already deploy MPLS, or plan to do so, can obtain these benefits from Cisco 6PE:

- There is minimal operational cost and risk and no impact on existing IPv4 and MPLS service, because only PE routers are upgraded. A Cisco 6PE router can be an existing PE router or a new router that is dedicated to IPv6 traffic.
- There is no impact on IPv6 CE routers. The ISP can connect to any CE device that runs static, IGP, or EGP ready-for-production services. An ISP can delegate IPv6 prefixes to customers.
- Nondisruptive IPv6 can be introduced into an existing MPLS service Cisco 6PE router at any time.

There are some drawbacks to using Cisco 6PE for integration:

- This option makes sense only when the network already runs MPLS.
- It requires knowledge of MPLS and BGP technologies.
- It requires dual stack and software upgrades on existing PE routers or deployment of dedicated Cisco 6PE routers.

https://t.me/learningnets



Follow these steps to deploy Cisco 6PE:

Step 1 Configure Cisco 6PE routers as dual-stack routers.

Step 2 Configure Cisco 6PE routers to be part of the IPv4 network and configure MPLS.

Step 3 Specify the interface from which locally generated packets take their source IPv6 addresses.

Step 4 Bind and advertise aggregate labels.

Example

In the network that is shown in the figure, two configuration tasks are required at the 6PE1 router to enable the Cisco 6PE feature.

The CE router, CE1, is configured to forward its IPv6 traffic to the 6PE1 router. The P1 router in the core of the network is assumed to be running MPLS, a label distribution protocol, an IPv4 IGP, and Cisco Express Forwarding or distributed Cisco Express Forwarding. P1 does not require any new configuration to enable the Cisco 6PE feature. New configuration tasks are also not required for the CE1 router.

The Cisco 6PE routers, 6PE1 and 6PE2, must be members of the core IPv4 network. The Cisco 6PE router interfaces that are attached to the core network must run MPLS, the same LDP, and the same IPv4 IGP as in the core network.

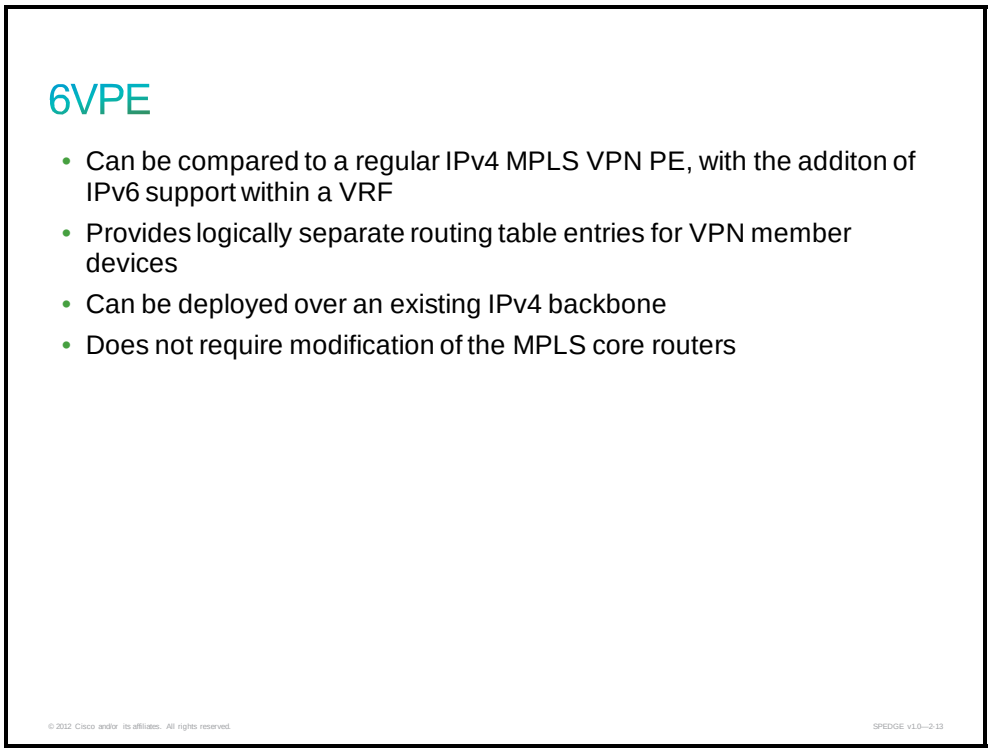
The 6PE routers must also be configured to be dual stack to run both IPv4 and IPv6.

https://t.me/learningnets

Restrictions

These restrictions apply when implementing the Cisco 6PE feature:

- Core MPLS routers support MPLS and IPv4 only, so they cannot forward or create any IPv6 Internet Control Message Protocol (ICMP) messages.
- Cisco 6PE does not provide load balancing between an MPLS path and an IPv6 path. If both are available, the MPLS path is always preferred. Load balancing between two MPLS paths is possible.
- When MP-IBGP multipath is enabled on the Cisco 6PE router, all labeled paths are installed in the forwarding table with MPLS information (label stack) when MPLS available. This functionality enables Cisco 6PE to perform load balancing.



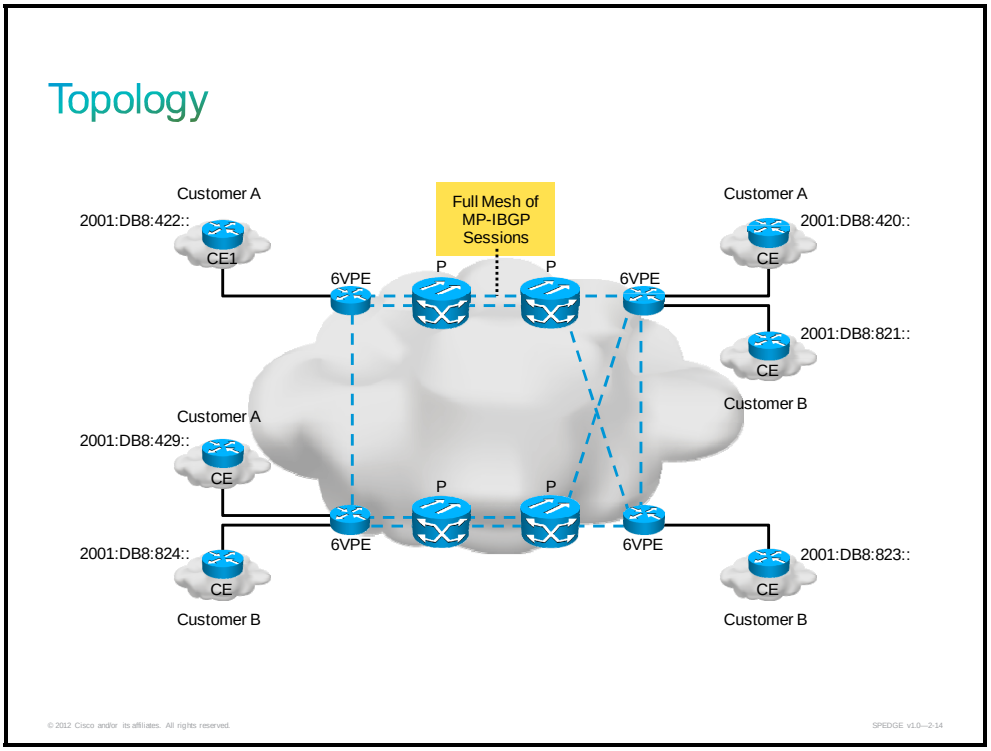
IPv6 MPLS VPNs are like IPv4 MPLS VPNs. Customer sites are connected over the shared MPLS backbone so that customers have reachability to their own VPN partner sites only. As with IPv4 MPLS VPNs, this outsourced intranet remains private and is seen by the customer as a closed network that runs on a dedicated network.

The inherent VPN and terminal equipment services that are available within an MPLS environment allow IPv6 networks to be combined into VPNs or extranets over an infrastructure that supports IPv4 VPNs and MPLS terminal equipment.

ISPs that offer MPLS VPNs for IPv4 can use Cisco 6VPE to add IPv6 services. This approach has these benefits:

- There is no modification on the MPLS core.
- Cisco 6VPE supports both IPv4 and IPv6 VPNs concurrently on the same interfaces.
- Configuration and operation of IPv6 VPNs are the same as for IPv4 VPNs.

https://t.me/learningnets



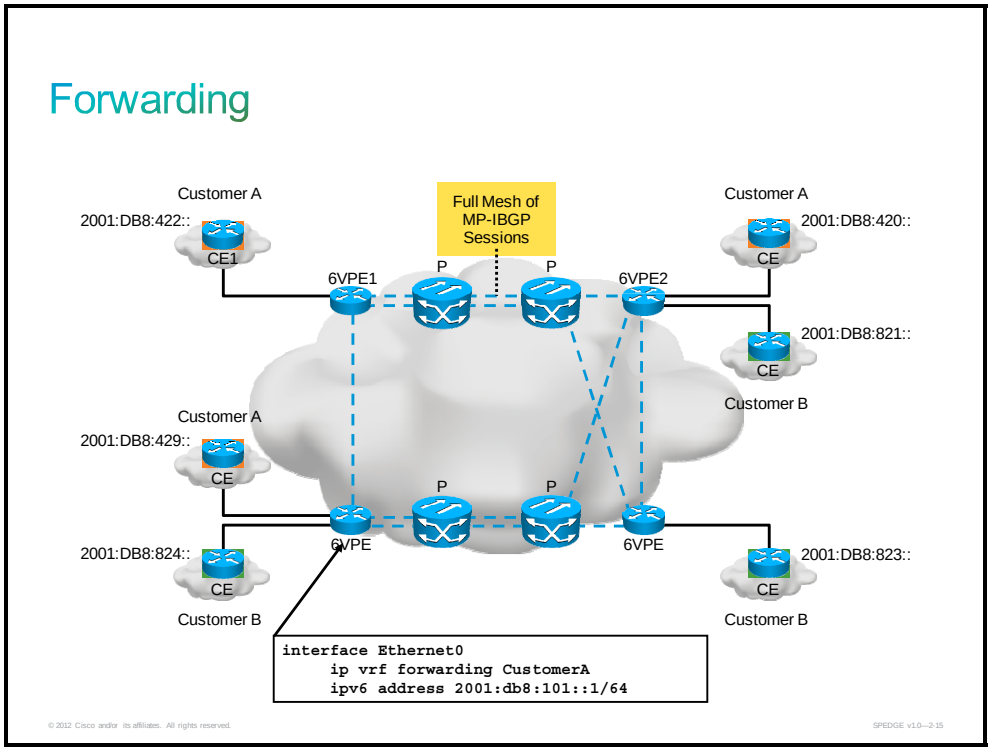
Service providers that offer MPLS IPv4 VPN services to their customers might look forward to adding IPv6 VPN services to their portfolios. A VPN is said to be an IPv6 VPN when a CE router turns on native IPv6 over an interface or subinterface to the PE router. Cisco 6VPE adds IPv6 VPN capability to Cisco 6PE and enables an ISP to deliver services that are similar to IPv4.

As with IPv4 VPN route distribution, BGP and its extensions are used to distribute routes from an IPv6 VPN site to all other Cisco 6VPE routers that connect to a site of the same IPv6 VPN. Cisco 6PE routers use virtual routing and forwarding (VRF) tables to separately maintain the reachability and forwarding information of each IPv6 VPN (as with IPv4 VPN).

The IETF standard defines a specification that lets a service provider use an MPLS-enabled IPv4 backbone to provide VPNs for its IPv6 customers. The standard method is based on VPNs as described in RFC 4364, *BGP/MPLS IP VPNs*. In a BGP/MPLS VPN, MPLS is used to forward packets over the backbone, and Multiprotocol Border Gateway Protocol (MP-BGP) is used to distribute VPN routes over the service provider backbone.

MP-BGP allows generic operations over both IPv4 and IPv6. For example, MP-BGP defines an IPv6 VPN address family and describes the route distribution and MPLS tunnel selection.

https://t.me/learningnets



Forwarding in Cisco 6VPE follows these steps:

1. The ingress PE router receives an IPv6 packet and attaches an MPLS label.

When a Cisco 6VPE router receives an IPv6 packet from CE1, the router looks up the IPv6 packet destination address in the VRF table. This address enables 6VPE1 to find a VPN IPv6 route. The VPN IPv6 router has an associated MPLS label and an associated BGP next hop. This MPLS label is imposed on the IPv6 packet.

2. The ingress PE router attaches a second label to the IPv6 packet.

6VPE1 pushes another label, which is bound by LDP IGP to the IPv4 address of the BGP next hop, to reach 6VPE2 through the MPLS cloud on the label stack of the labeled IPv6 VPN packet. This top-most imposed label corresponds to the LSP that starts on 6VPE1 and ends on 6VPE2. The bottom label is bound to the IPv6 VPN prefix via BGP. All the P routers in the backbone network switch the VPN packet, based only on the top label in the stack, which points toward the 6VPE2 router. Because of the normal MPLS forwarding rules, the P routers never look beyond the first label and are thus completely unaware of the second label or the IPv6 VPN packet that is carried across the backbone network.

3. The egress router receives the labeled IPv6 packet, removes the labels, and forwards the packet to the CE router.

The egress PE router, 6VPE2, receives the labeled IPv6 VPN packet, drops the first label, and performs a lookup on the second label, which uniquely identifies the target VRF table and sometimes even the outgoing interface on 6VPE2. A lookup is performed in the target VRF table, and the IPv6 packet is sent toward the proper CE router in the IPv6 domain or site.

https://t.me/learningnets

Configuring 6VPE

The configuration steps for 6VPE on PE routers are as follows:

- Enable IPv6 routing and IPv6 Cisco Express Forwarding.
- Configure MP-BGP for VPN route exchange.
- Configure the customer IPv6 VRF tables.
- Configure the customer VRF interfaces.
- Configure the customer VRF (PE-CE) IPv6 routing protocols or static IPv6 routes.
- Redistribute the CE-PE routing protocol VRF routes into MP-BGP.

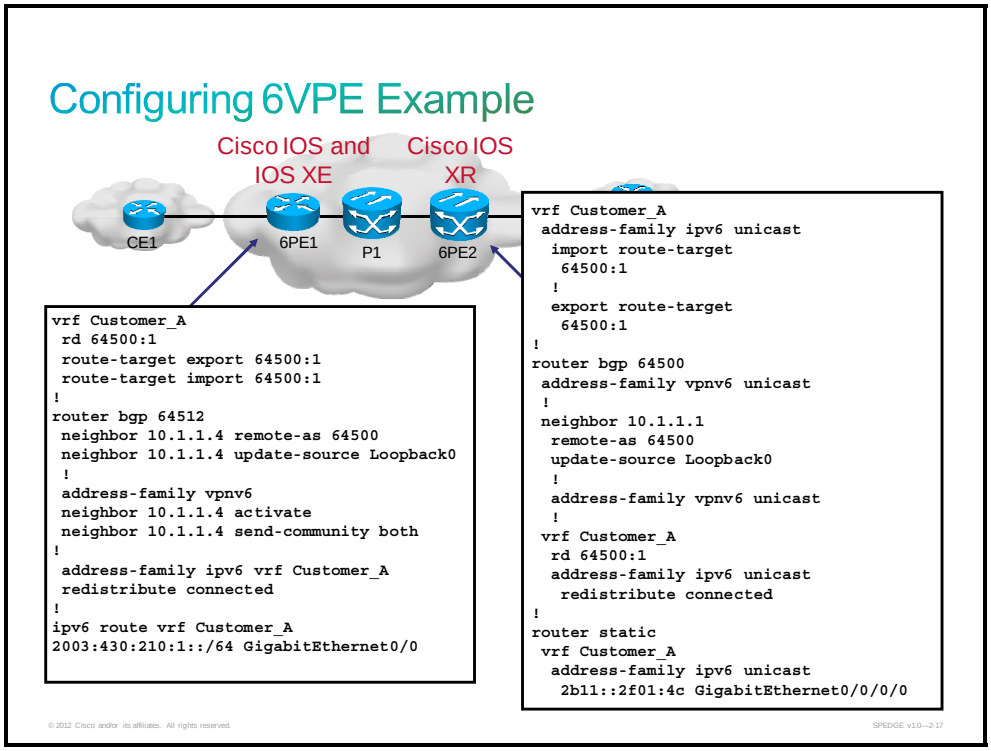
© 2012 Cisco and/or its affiliates. All rights reserved.

To successfully implement 6VPE, you must apply all the steps that the figure shows.

This task assumes that these steps have already been completed:

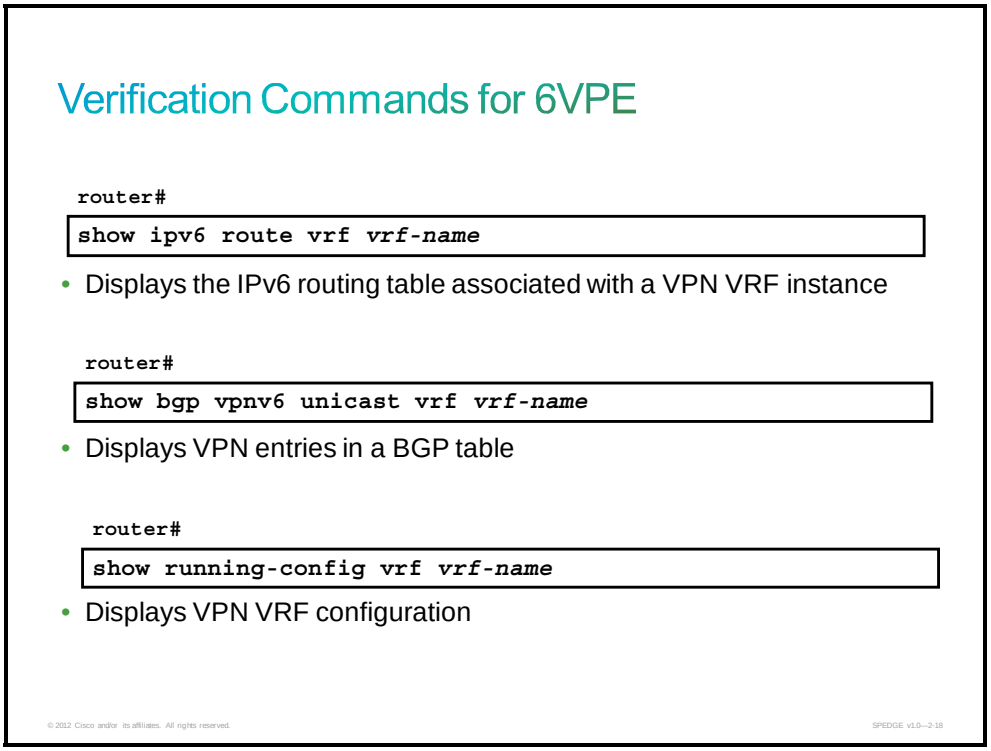
- A loopback interface is configured.
- LDP is configured.
- MPLS is enabled on interfaces.
- The backbone IGP (OSPF or IS-IS) is configured.

https://t.me/learningnets



The figure shows a sample configuration of PE routers. Router 6PE1 is running Cisco IOS and IOS XE Software, while 6PE2 is running Cisco IOS XR Software. Both PE routers have an MP-BGP session for route exchange and a configured VPN (Customer_A).

https://t.me/learningnets



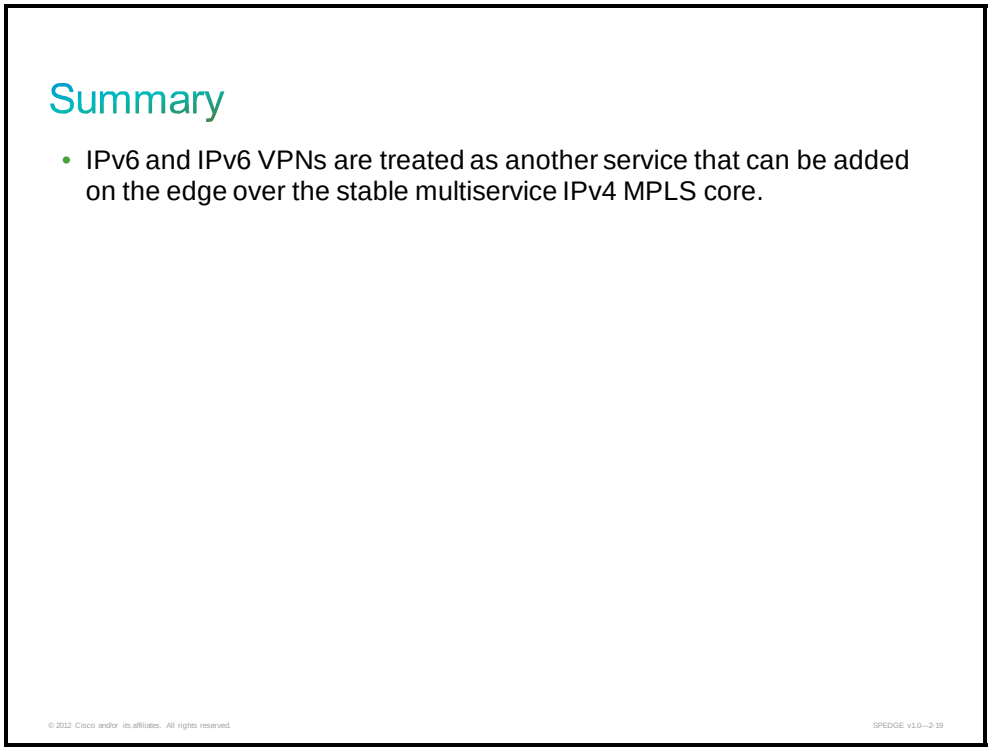
The figure shows commonly used commands for configuration verification.

```
6PE#show ipv6 route vrf VPN_A
IPv6 Routing Table - VPN_A - 6 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, R - RIP, I1 - ISIS L1, I2 - ISIS L2
       IA - ISIS interarea, IS - ISIS summary
       O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
C   2003:410:210:1::/64 [0/0]
     via Serial2/0, directly connected
L   2003:410:210:1:2D0:63FF:FE54:7000/128 [0/0]
     via Serial2/0, receive
B   2003:420:210:1::/64 [200/0] (line 1)
     via 10.1.1.4%Default-IP-Routing-Table, indirectly connected
S   2003:430:210:1::/64 [1/0] (line 2)
     via Serial2/0, directly connected
B   2003:440:210:1::/64 [200/0] (line 3)
     via 10.1.1.4%Default-IP-Routing-Table, indirectly connected
L   FF00::/8 [0/0]
     via Null0, receive
```


https://t.me/learningnets

Summary

This topic summarizes the primary point that was discussed in this lesson.

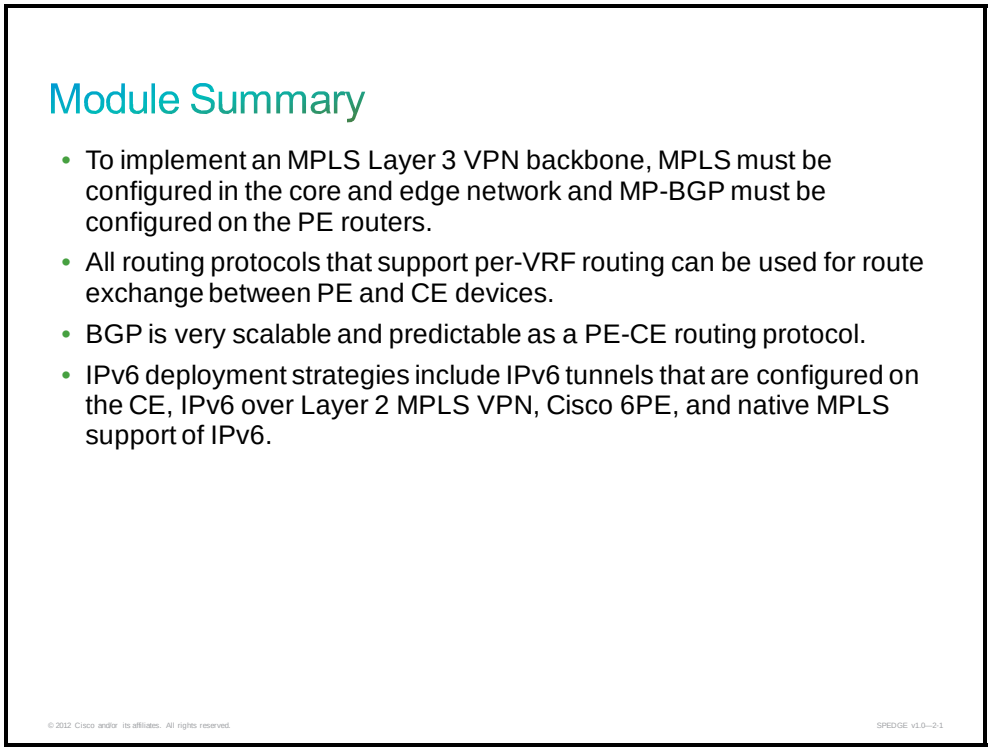


https://t.me/learningnets

https://t.me/learningnets

Module Summary

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



A Multiprotocol Label Switching (MPLS) VPN implementation involves virtual routing and forwarding (VRF) tables, the interaction between customer edge (CE)-to-provider (P) routing protocols, and Multiprotocol Border Gateway Protocol (MP-BGP) in the service provider backbone.

References

For additional information, refer to these resources:

- Cisco Systems, Inc. “Configuring MPLS Layer 3 VPNs.” http://www.cisco.com/en/US/docs/ios/mpls/configuration/guide/mp_cfg_layer3_vpn.html.
- Cisco Systems, Inc. “Implementing MPLS Layer 3 VPNs on Cisco IOS XR Software.” http://www.cisco.com/en/US/docs/ios_xr_sw/iosxr_r3.7/mpls/configuration/guide/gc37v3.html.
- Cisco Systems, Inc. The “MPLS VPN Implementation” module in the *Implementing Cisco MPLS* (MPLS) v2.3 course.
- Cisco Systems, Inc. The “Identifying IPv6 Service Provider Deployment” module in the *IPv6 Fundamentals, Design, and Deployment* (IP6FD) course.

https://t.me/learningnets

https://t.me/learningnets

Module Self-Check

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

Q1) In an MPLS VPN implementation, what is a VRF? (Source: Implementing MPLS Layer 3 VPN Backbones)

A) the routing and forwarding instance for all sites belonging to a single customer
B) the routing and forwarding instance for all sites belonging to a single customer location
C) the routing and forwarding instance for all sites using a common routing protocol
D) the routing and forwarding instance for a set of sites with identical connectivity requirements

Q2) Which two protocols are VPN-aware? (Choose two.) (Source: Implementing MPLS Layer 3 VPN Backbones)

A) RIPv2
B) IS-IS
C) ODR
D) EIGRP
E) RIPv1

Q3) VRFs are assigned to an interface. (Source: Implementing MPLS Layer 3 VPN Backbones)

A) true
B) false

Q4) Which command do you use to create a VRF named “VPNA” on Cisco IOS XR devices? (Source: Implementing MPLS Layer 3 VPN Backbones)

A) **ip vrf VPNA**
B) **vrf VPNA**
C) **ip vrf forwarding VPNA**
D) **vrf forwarding VPNA**

Q5) Which command do you use to associate interface Gigabit Ethernet 0/0/0/0 with a VRF named “VPNA” on Cisco IOS XR devices? (Source: Implementing MPLS Layer 3 VPN Backbones)

A) **ip vrf VPNA**
B) **ip vrf VPNA interface GigabitEthernet0/0/0/0**
C) **ip vrf forwarding VPNA**
D) **vrf VPNA**
E) **vrf VPNA interface GigabitEthernet0/0/0/0**

Q6) What happens to the existing IP address of an interface when you associate the interface with a VRF? (Source: Implementing MPLS Layer 3 VPN Backbones)

A) It remains unchanged.
B) It is removed from the interface.
C) It is changed to the Loopback 0 address.
D) It is moved under the VRF configuration.

https://t.me/learningnets

Q7) Which command configures the redistribution of static VRF routes between PE routers that are running Cisco IOS XR Software? (Source: Connecting Customers Using Simple Routing Protocols)

A) RP/0/RP0/CPU0:router(config)# **redistribute static**
B) RP/0/RP0/CPU0:router(config-router)# **redistribute static**
C) RP/0/RP0/CPU0:router(config-bgp)# **redistribute static**
D) RP/0/RP0/CPU0:router(config-bgp-vrf)# **redistribute static**
E) RP/0/RP0/CPU0:router(config-bgp-vrf-af)# **redistribute static**

Q8) How is each routing context implemented in OSPF? (Source: Connecting Customers Using BGP or OSPF)

A) by redistributing into MP-BGP
B) as a separate routing process
C) by assigning it to an interface
D) as a separate isolated instance of the same routing protocol

Q9) In the SOO extended community attribute in MP-BGP, what does “SOO” stand for ? (Source: Connecting Customers Using BGP or OSPF)

A) Site of Origin
B) Sites of Origin
C) State of Origin
D) Source of Origin

Q10) Which command do you use to display the contents of the LFIB table on a PE router running Cisco IOS XR Software? (Source: Connecting Customers Using BGP or OSPF)

A) **show mpls forwarding-table**
B) **show mpls forwarding-table vrf**
C) **show mpls forwarding**
D) **show mpls forwarding vrf**

Q11) No tunnel or circuit configuration is required when using Cisco 6PE for IPv6 integration. (Source: Deploying IPV6 and MPLS)

A) true
B) false

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Module Self-Check Answer Key

Q1) D

Q2) A, D

Q3) B

Q4) B

Q5) D

Q6) B

Q7) E

Q8) B

Q9) A

Q10) C

Q11) A

https://t.me/learningnets

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Module 3

Complex MPLS Layer 3 VPNs

Overview

This module discusses some advanced configuration features that can help increase the stability of the Multiprotocol Label Switching (MPLS) VPN backbone. The first lesson discusses various MPLS VPN features that a service provider might use to help meet service requirements, and looks at various types of VPN solutions and topologies.

Integrating Internet access with an MPLS VPN solution is one of the most common service provider business requirements. The second lesson provides a good understanding of underlying design issues, several potential design scenarios, and some sample configurations. Various topologies and implementation methods are discussed, along with ways to separate Internet access from VPN services.

Deployments of MPLS have become routine in large-scale global networks, which demand solutions to complex business and network problems. The third lesson describes the two primary components of the Cisco IOS MPLS Inter-Domain Solution: interautonomous system (inter-AS) and Carrier Supporting Carrier (CSC).

Module Objectives

Upon completing this module, you will be able to describe and implement the most important complex Layer 3 VPN features. You will be able to meet these objectives:

- Describe the requirements, usage, and solutions that are associated with overlapping, central services, and managed CE router service VPNs
- Describe common customer Internet connectivity scenarios and identify design models for combining Internet access with MPLS Layer 3 VPN services
- Introduce CSC and inter-AS features with different implementation options

https://t.me/learningnets

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 1

Implementing Complex MPLS Layer 3 VPNs

Overview

This lesson discusses some advanced configuration features that can help increase the stability of the Multiprotocol Label Switching (MPLS) VPN backbone. The lesson also discusses various MPLS VPN features that a service provider might use to help meet service requirements, and looks at various types of VPN solutions and topologies.

Objectives

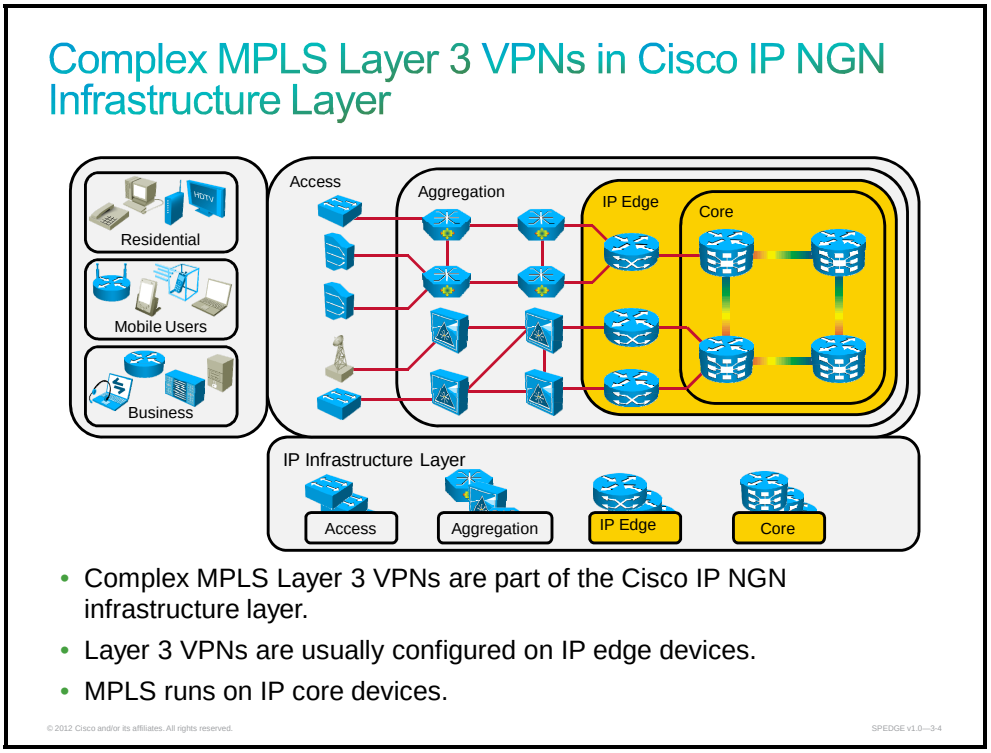
Upon completing this lesson, you will be able to describe the requirements, usage, and solutions associated with overlapping, central services, and managed CE router service VPNs. You will be able to meet these objectives:

- Describe overlapping VPNs
- Describe central service VPNs and advanced VRF features
- Describe managed CE router service

https://t.me/learningnets

Overlapping VPNs

This topic describes how a service provider can configure overlapping VPNs. Overlapping VPNs are usually used to connect parts of two separate VPNs.

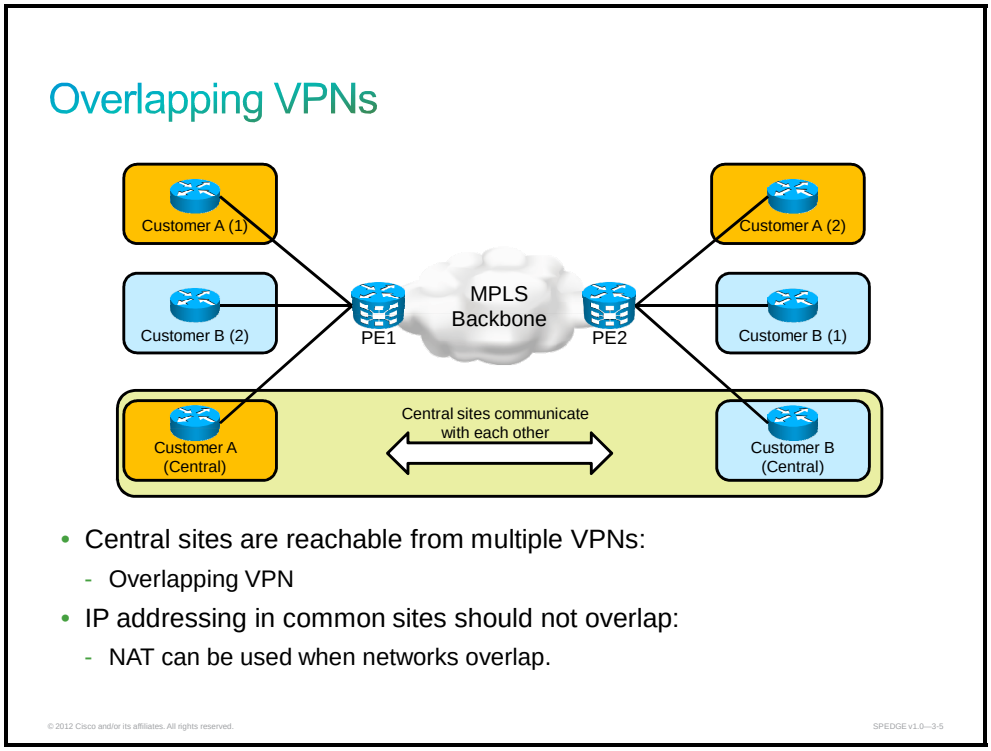


Complex MPLS Layer 3 VPNs include the following traits:

- Complex MPLS Layer 3 VPNs are part of the Cisco IP Next-Generation Network (NGN) infrastructure layer
- Layer 3 VPNs are usually configured on IP edge devices
- MPLS runs on IP core devices

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets



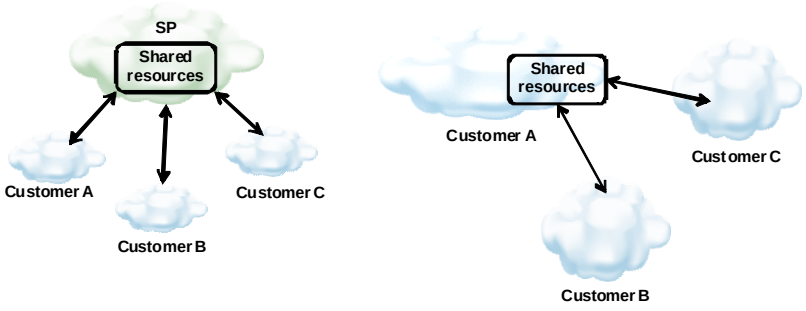
When two VPN customers want to share information, they might decide to interconnect their central sites. To achieve this interconnection, two simple VPNs are created, each containing a customer central site and its remote sites. Then a third VPN, which partially overlaps with the customer VPNs but connects only their central sites, is created. The central sites can talk to each other. Each central site can also talk to the remote sites in its simple VPN, but not to the remote sites belonging to the other customer simple VPN. The addresses used in the central sites, however, must be unique in both VPNs.

Another option is to use dual Network Address Translation (NAT) with a registered address to be imported and exported between the two central sites.

https://t.me/learningnets

Typical Overlapping VPN Usages

- At least one customer site needs to be reachable by sites in different VPNs:
  - A service provider may provide services to many customers.
  - Some service provider customers may want connectivity to one of their partners through the MPLS network.
  - Limit visibility between different departments in an organization.



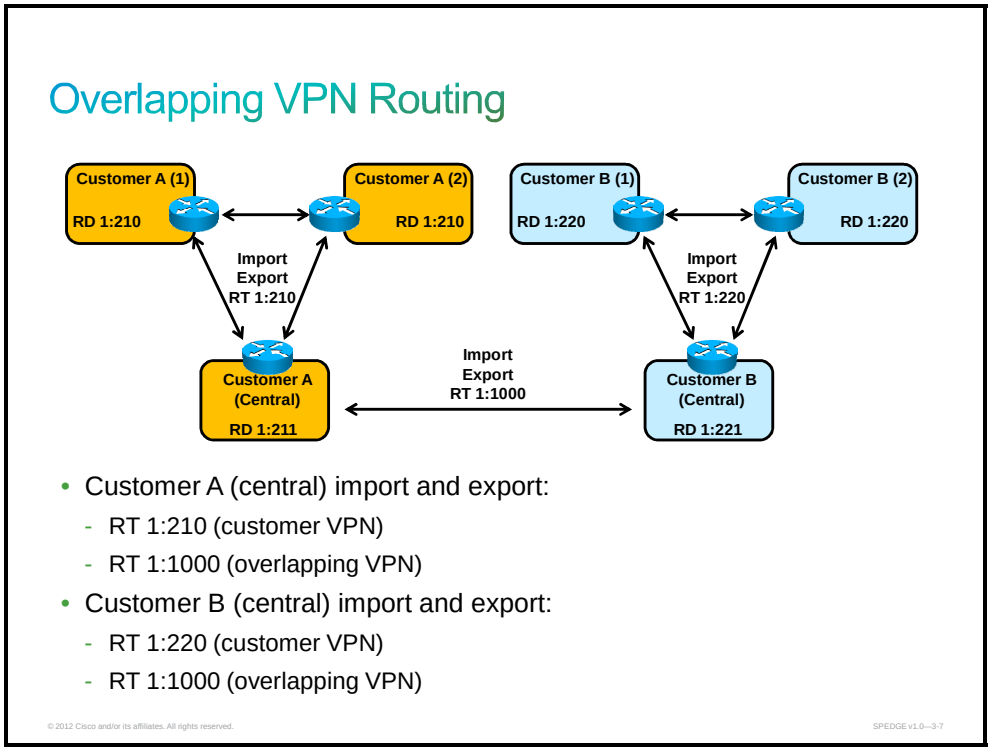
© 2012 Cisco and/or its affiliates. All rights reserved.

Typical uses for overlapping VPNs include the following:

- A service provider may provide services to its customers.
- Companies that use MPLS VPNs to implement both intranet and extranet services might use overlapping VPNs. In this scenario, each company participating in the extranet VPN would probably deploy a security mechanism on its customer edge (CE) routers to prevent other companies that are participating in the VPN from gaining access to other sites in its VPN.
- A security-conscious company might decide to limit visibility between different departments in the organization. Overlapping VPNs might be used as a solution.

Note Security issues might force an enterprise network to be migrated to an MPLS VPN even if it is not using MPLS VPN services from a service provider.

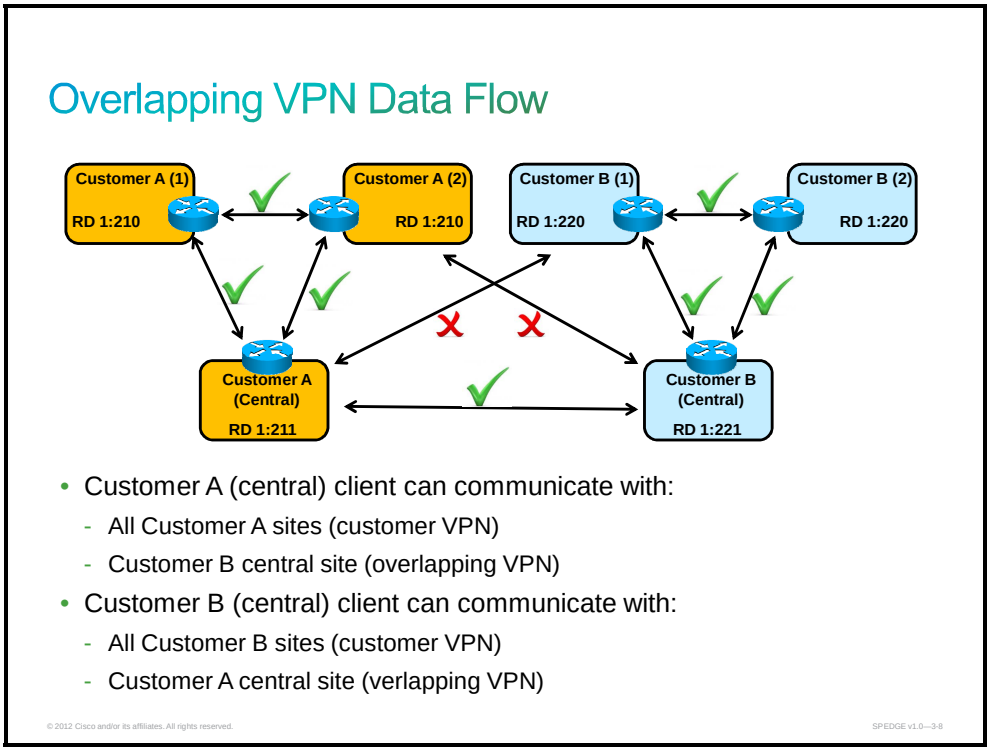
https://t.me/learningnets



The figure shows how to implement overlapping VPNs between two customer central sites. Only the configuration on the customer central site PE router has to be changed:

- New virtual routing and forwarding (VRF) is created for the customer central site.
- A new router distinguisher (RD) is configured for customer central site.
- For the Customer A central site, the following occurs:
  - Import and export routes with RT 1:210 (customer routes)
  - Import and export routes with RT 1:1000 (overlapping routes)
- For the Customer B central site, the following occurs:
  - Import and export routes with RT 1:220 (customer routes)
  - Import and export routes with RT 1:1000 (overlapping routes)

https://t.me/learningnets

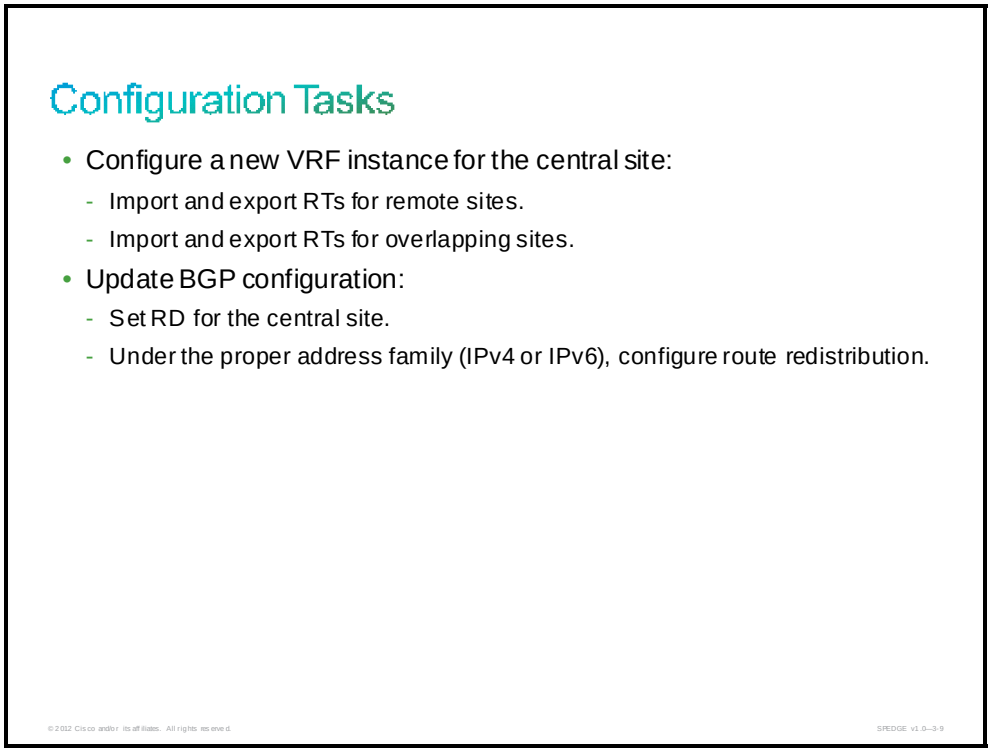


Because sites belonging to different VPNs do not share routing information, they cannot talk to each other. The figure shows overlapping VPN data flow:

- The simple VPN for Customer A contains routes that originate from the following:
  - A-Central site
  - A remote sites
- The simple VPN for Customer B contains routes that originate from the following:
  - B-Central site
  - B remote sites
- The overlapping VPN contains routes that originate from the following:
  - A-Central site
  - B-Central site
- All Customer A sites can communicate with each other.
- All Customer B sites can communicate with each other.
- A-Central and B-Central can communicate with each other.
- The customer A remote site cannot communicate with the customer B remote sites.
- The customer A central site cannot communicate with the customer B remote sites.

Note If a site participating in more than one VPN is propagating a default route to other sites, it can attract traffic from those sites and start acting as a transit site between VPNs, enabling sites that were not supposed to communicate to establish two-way communication.

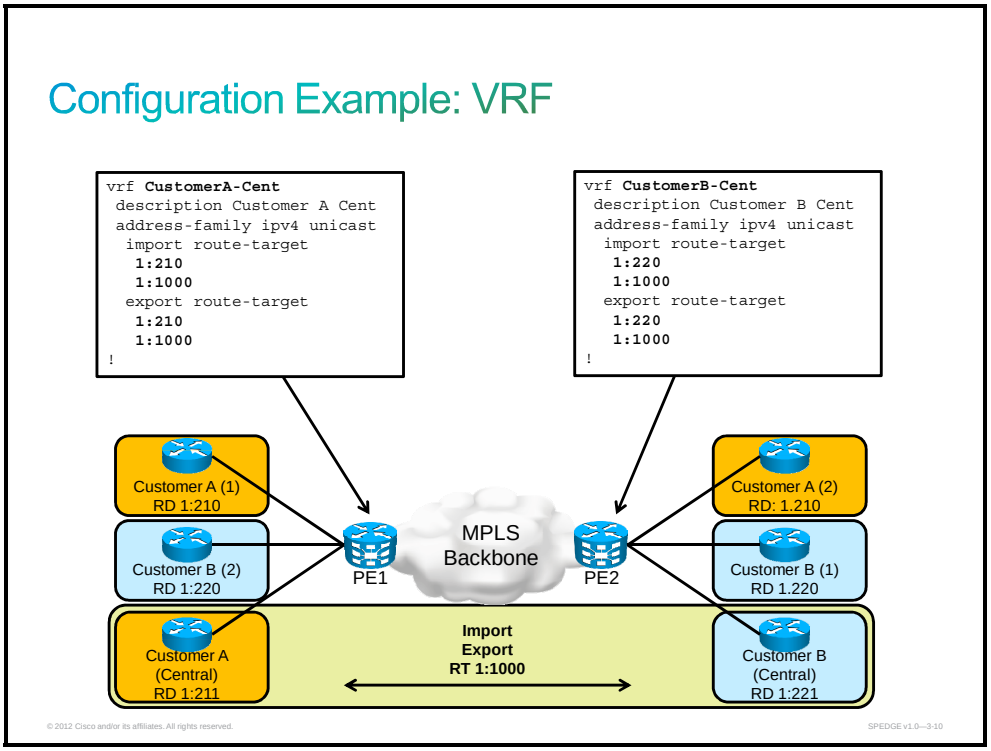
https://t.me/learningnets



To configure overlapping VPNs, the administrator has to do the following:

- Configure a new VRF instance for the central site
  - Import and export route targets (RTs) for remote sites
  - Import and export RTs for the overlapping site
- Update BGP configuration
  - Set RDs for the central site
  - Under the proper address family (IPv4 or IPv6), configure route redistribution

https://t.me/learningnets

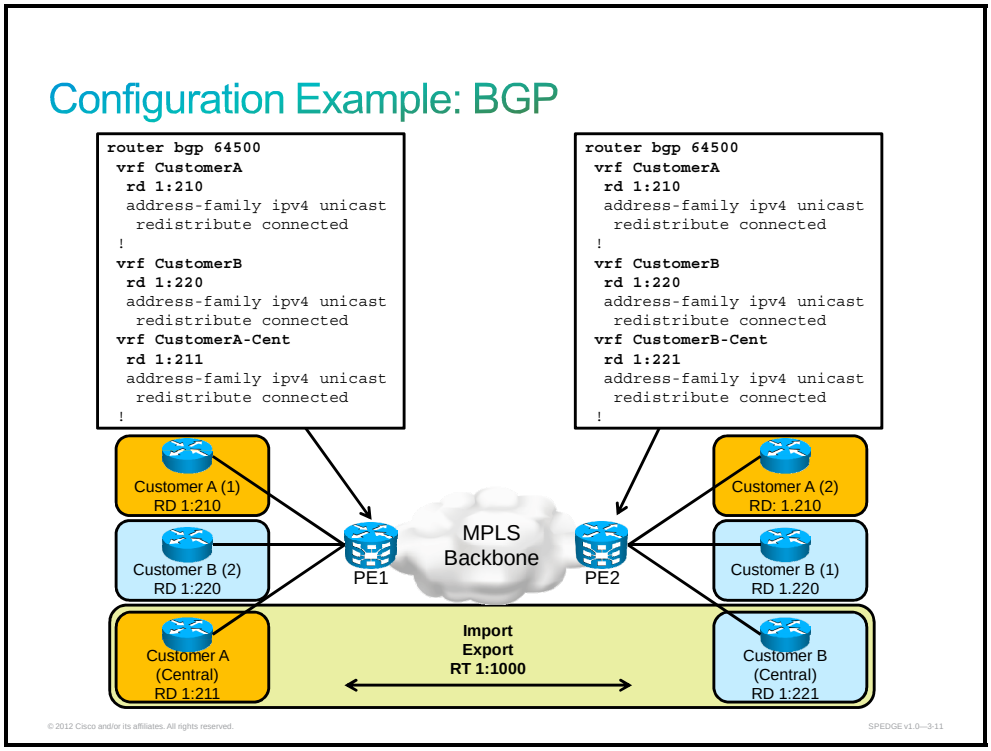


The figure shows a Cisco IOS XR configuration example for configuring VRF instances.

Provider edge (PE) configuration is as follows:

```
vrf CustomerA-Cent
 description Customer A Cent
 address-family ipv4 unicast
  import route-target
   1:210
   1:1000
  export route-target
   1:210
   1:1000
!
vrf CustomerB-Cent
 description Customer B Cent
 address-family ipv4 unicast
  import route-target
   1:220
   1:1000
  export route-target
   1:220
   1:1000
!
vrf CustomerA
 description Customer A
 address-family ipv4 unicast
  import route-target
   1:210
  export route-target
   1:210
!
vrf CustomerB
 description Customer B
 address-family ipv4 unicast
  import route-target
   1:220
  export route-target
   1:220
!
```


https://t.me/learningnets



The figure shows a Cisco IOS XR configuration example for configuring the Border Gateway Protocol (BGP) process.

PE configuration is as follows:

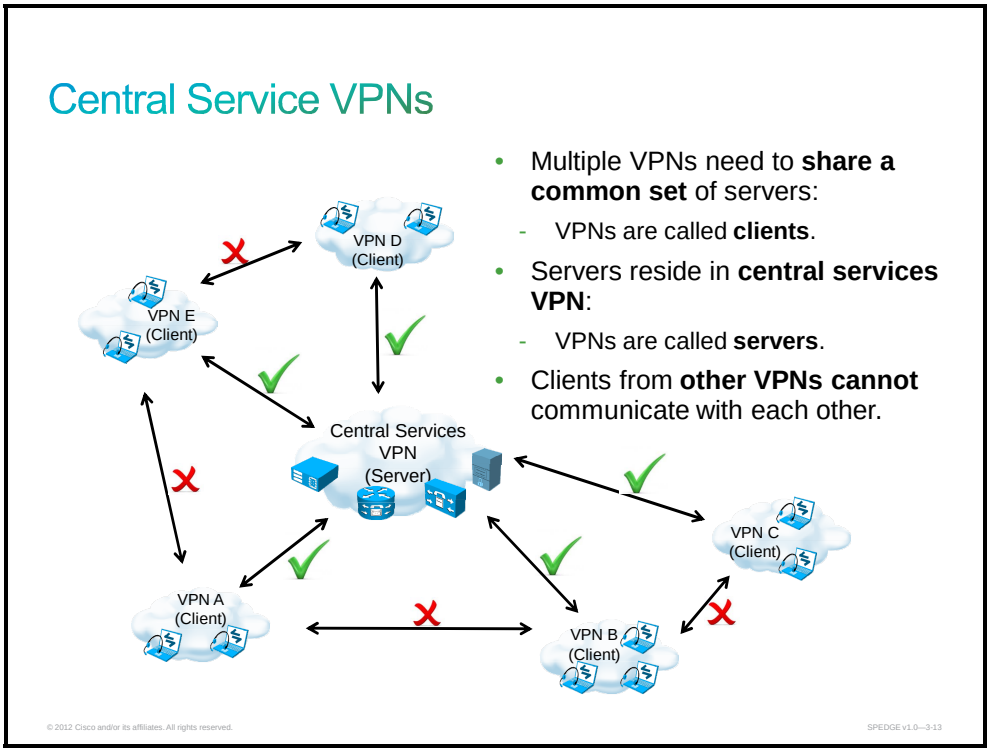
```
router bgp 64500
 vrf CustomerA
  rd 1:210
  address-family ipv4 unicast
   redistribute connected
 !
 vrf CustomerB
  rd 1:220
  address-family ipv4 unicast
   redistribute connected
!
 vrf CustomerA-Cent
  rd 1:211
  address-family ipv4 unicast
   redistribute connected
!
 vrf CustomerB-Cent
  rd 1:221
  address-family ipv4 unicast
   redistribute connected

 !
```


https://t.me/learningnets

Central Service VPNs

This topic describes central service VPNs and advanced VRF features.



A central service VPN is a topology with these characteristics:

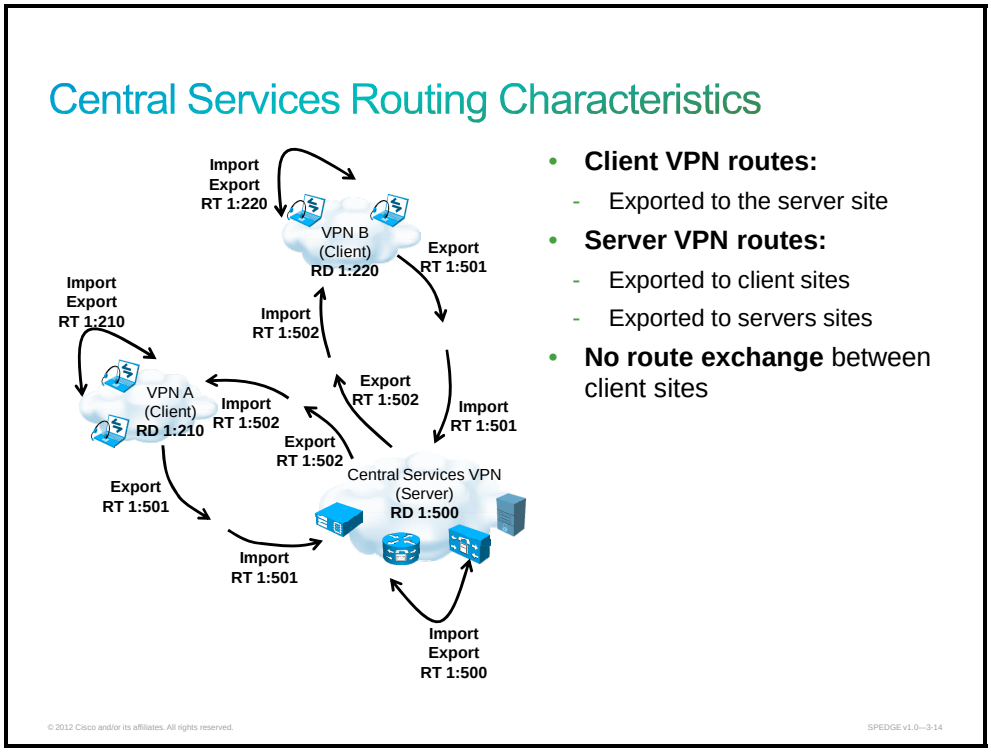
- Some sites (“server sites”) can communicate with all other sites.
- All the other sites (“client sites”) can communicate only with the server sites.

This topology can be used in these situations:

- The service provider offers services to all customers by allowing them access to a common VPN.
- Two (or more) companies want to exchange information by sharing a common set of servers.

A security-conscious company separates its departments and allows them access only to common servers.

https://t.me/learningnets



There is a specific routing model used to implement a central services VPN.

The figure illustrates the MPLS VPN routing model that is used to implement a central services VPN:

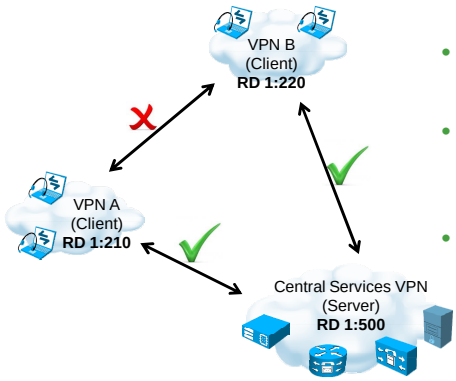
- Client site VPN routes are exported to the server site VPN.
- Server site VPN routes are exported to client site VPNs
- Server site VPN routes are also exchanged between other server sites.
- There should be no route exchange and connectivity between client sites.

The figure shows a central services VPN (RD 1:500) with a set of common services. On client site VPNs, routes are exported with RT 1:501 and imported with RT 1:502.

On the server site VPN, routes are exported with RT 1:502 and imported with RT 1:501.

https://t.me/learningnets

Central Services Data Flow



- Clients can talk to servers:
  - Client VRF contains server routes.
- Servers can talk to clients:
  - Server VRF contains client routes.
- Clients cannot communicate:
  - Client VRFs do not contain routes from other clients;
- Make sure that there is no client-to-client leakage across server sites.

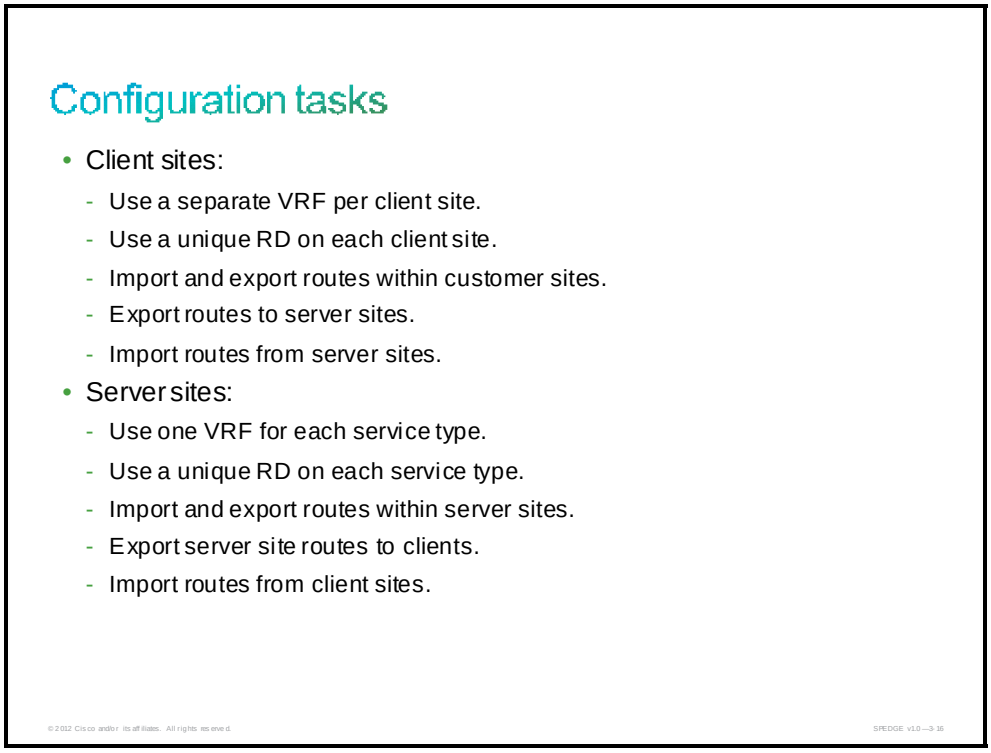
© 2012 Cisco and/or its affiliates. All rights reserved.

In the central services VPN topology, the client VRF contains only routes from the client site and routes from the server sites. This setup precludes the client sites from communicating with other client sites.

A server VRF in this topology contains routes from the site or sites attached to the VRF and also routes from all other client and server sites. Hosts in server sites can therefore communicate with hosts in all other sites.

Note	If the central site is propagating a default route to other sites, it can result in client sites seeing each other through the CE router in the central site.

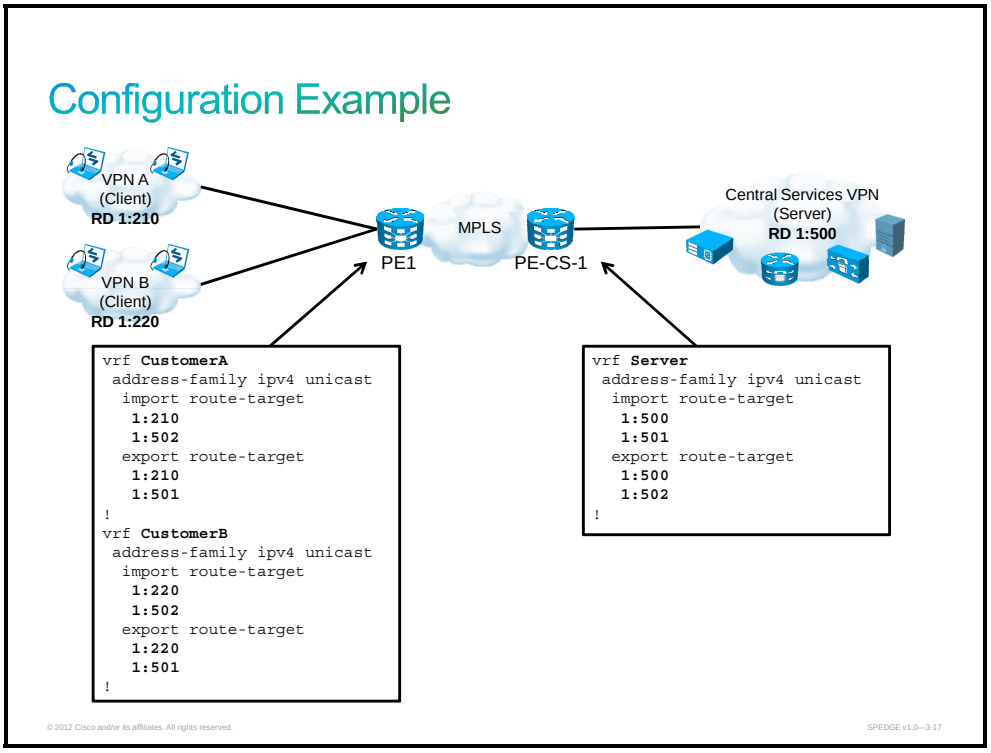
https://t.me/learningnets



To configure a central services VPN, you need to address these requirements:

- Client sites
  - Use a separate VRF per client site
  - Use a unique RD on each client site
  - Import and export routes within customer sites
  - Export routes to server sites
  - Import routes from server sites
- Server sites
  - Use one VRF for each service type
  - Use a unique RD on each service type
  - Import and export routes within the server site
  - Export server site routes to clients
  - Import routes from client sites

https://t.me/learningnets



The configuration example in the figure shows how to configure a central services VPN in a Cisco IP NGN service provider network.

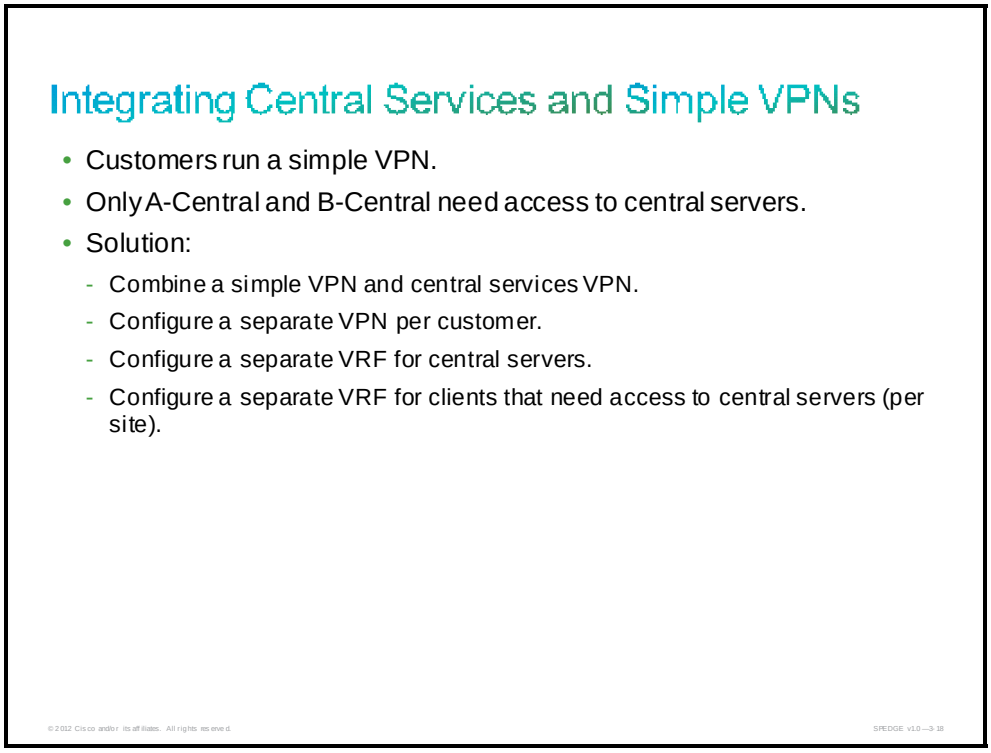
On PE1, two VRF instances are configured for Customer A and Customer B. On the PE-CS-1 router, the VRF server is configured.

Router PE-CS-1 exports routes with RT 1:502 and imports client routes with RT 1:501.

Customer VRF exports routes with RT 1:501 and imports server routes with RT 1:502.

 © 2012 Cisco Systems, Inc.

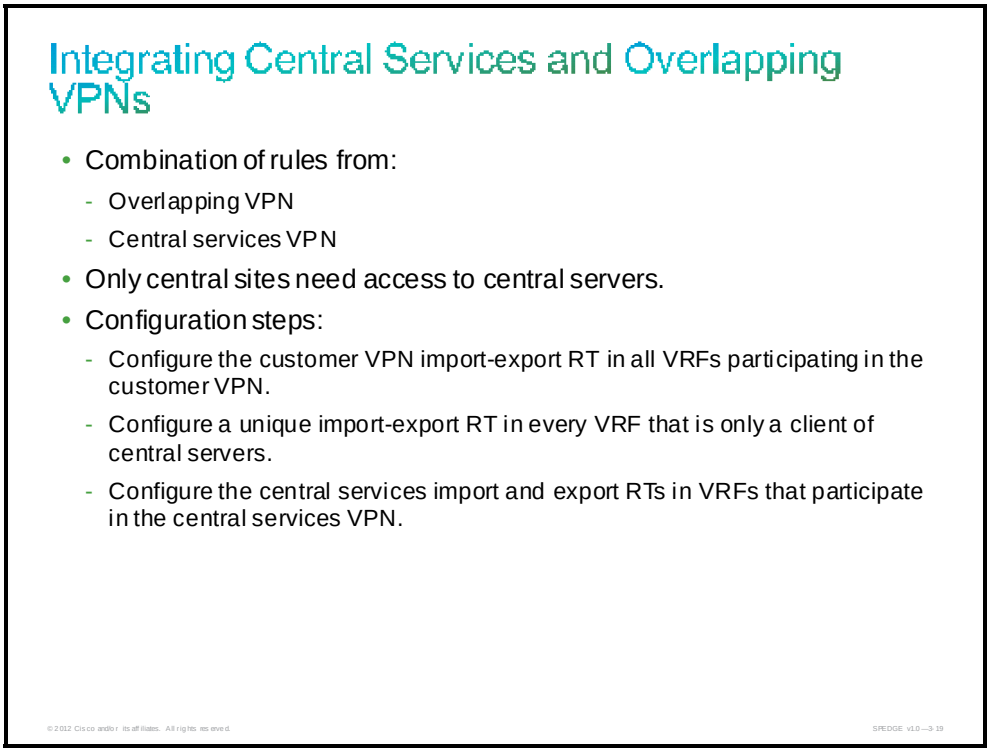
https://t.me/learningnets



In this design, some of the customer sites need access to the central server. All other sites just need optimal intra-VPN access. The design is consequently a mixture of simple VPN topology and central services VPN topology.

When integrating a central services VPN with a simple VPN, you need one VRF per VPN for sites that have access to other sites in the customer VPN but that have no access to the central services VPN. You need one VRF per VPN for sites that have access to the central services VPN. Finally, you need one VRF for the central services VPN.

https://t.me/learningnets

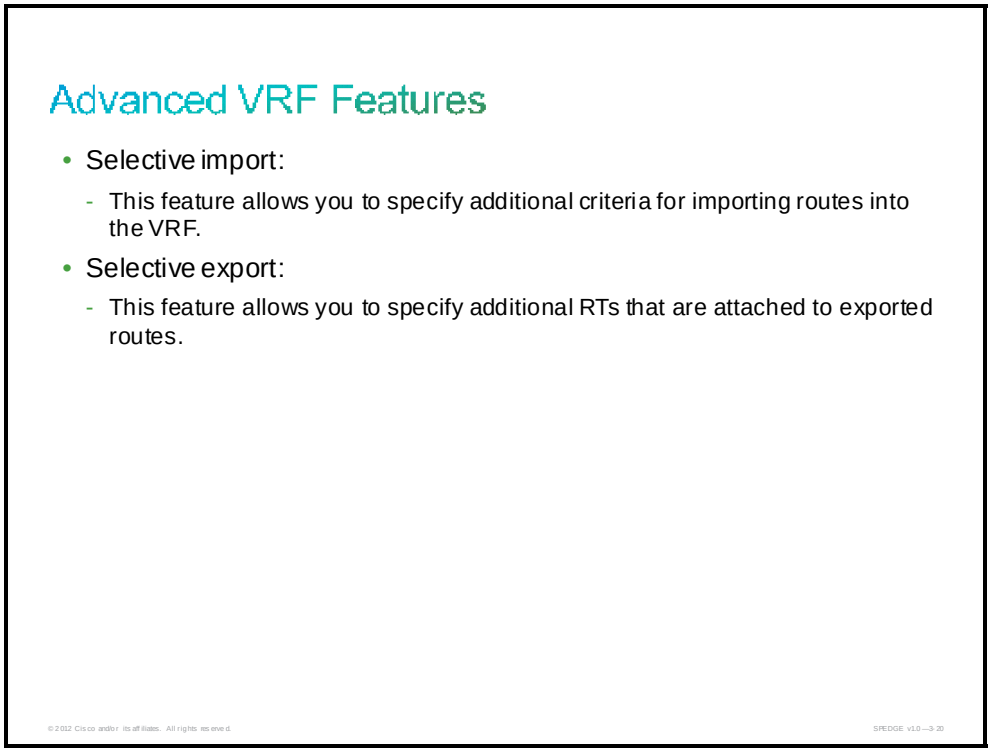


To integrate central services and overlapping VPNs, you have to combine rules from overlapping VPNs and central services VPNs.

The configuration steps are as follows:

- Configure the customer VPN import-export RT in all VRFs that are participating in the customer VPN.
- Configure a unique import-export RT in every VRF that is only a client of central servers.
- Configure the central services import and export RTs in VRFs that participate in the central services VPN.

https://t.me/learningnets



These advanced VRF features allow you to deploy advanced MPLS VPN topologies or increase the stability of the MPLS VPN backbone:

- The selective import feature allows you to select routes to be imported into a VRF based on criteria other than the RT of the VRF.
- The selective export feature allows you to attach specific RTs to a subset of routes that are exported from a VRF. By default, the same RTs get attached to all exported routes.

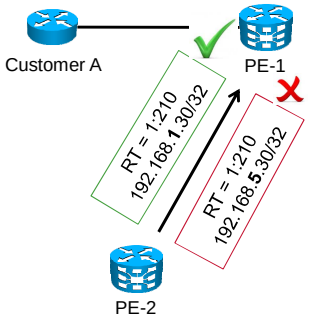
Note	The VRF route limit is also an advanced VRF feature on some platforms that allows you to limit the number of routes that the customer—or other PE routers—can insert in the VRF. This feature prevents undesirable consequences such as configuration errors or denial-of-service (DoS) attacks.

https://t.me/learningnets

Selective VRF Import

- VRF import criteria are more specific than just the match in RT:
  - Import only routes with specific BGP attributes
  - Import routes with specific prefixes or subnet masks
- Route policy is used to make the route import selection more specific.
- Use the **import route-policy** *<name>* command in VRF configuration submode.

```
PE-1#
vrf CustomerA
 address-family ipv4 unicast
  import route-policy CustA-Policy
  import route-target
   1:210
  !
  export route-target
   1:210
!
route-policy CustA-Policy
  if destination in (192.168.1.0/24) then
    pass
  endif
end-policy
```



© 2012 Cisco and/or its affiliates. All rights reserved.

Selective route import into a VRF allows you to narrow the route import criteria. Selective route import uses a route policy that can filter the routes selected by the RT import filter. The routes imported into a VRF are Border Gateway Protocol (BGP) routes, so you can use match conditions in a route policy to match any BGP attribute of a route. These attributes include communities, local preference, multi-exit discriminator (MED), autonomous system (AS) path, and so on.

The **import route-policy** command is combined with the RT import filter. A route must pass the RT import filter first and then the import route policy. The necessary conditions for a route to be imported into a VRF are as follows:

- At least one of the RTs attached to the route matches one of the import RTs configured in the VRF.
- The route is permitted by the import route policy.

The figure shows an example in which an import route policy is used to match the IPv4 portion of incoming VPNv4 routes and to import into the VRF only routes matching a certain prefix.

A configuration similar to this one could be used to accomplish the following:

- Deploy advanced MPLS VPN topologies (for example, a managed router services topology)
- Increase the security of an extranet VPN by allowing only predefined subnetworks to be inserted into a VRF, thus preventing an extranet site from inserting unapproved subnetworks into the extranet

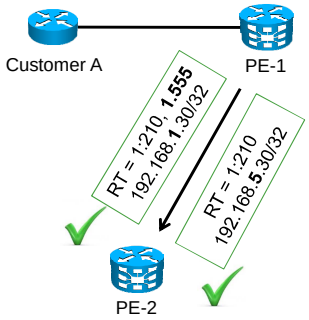
Note	A similar function is usually not needed in an intranet scenario because all customer routers in an intranet are usually under common administration.

https://t.me/learningnets

Selective VRF Export

- Routes from a VRF might have to be exported with different RTs:
  - Export management routes with particular RTs.
- An export route policy is used to set extended community RTs.

```
PE-1#
vrf CustomerA
 address-family ipv4 unicast
  import route-target
   1:210
   !
  export route-policy ExportPol
  export route-target
   1:210
!
route-policy ExportPol
 if destination in (192.168.1.0/24) then
    set extcommunity rt 1:555 additive
  else
    pass
  endif
end-policy
```



© 2012 Cisco and/or its affiliates. All rights reserved.

Some advanced MPLS VPN topologies are easiest to implement if you can attach various RTs to routes exported from the same VRF. This capability allows only a subset of the routes exported from a VRF to be imported into another VRF. Most services in which customer routers need to connect to a common server (for example, network management stations, voice gateways, and application servers) fall into this category.

The **export route-policy** command provides exactly this functionality. A route policy can be specified for each VRF to attach additional RTs to routes exported from that VRF. The **export route-policy** command performs only the attachment of RTs. It does not perform any filtering functions.

Attributes attached to a route with an export route policy are combined with the export RT attributes. If you specify export RTs in a VRF and set RTs with an export route policy, all specified RTs will be attached to the exported route.

Note The export route policy provides functionality that is almost identical to that of the import route map, but applied to a different VRF. Any requirement that can be implemented with an export route policy can also be implemented with an import route policy. However, the implementation of export maps can be more complicated and difficult to manage.

In the figure, the configuration is implemented with an export route policy.

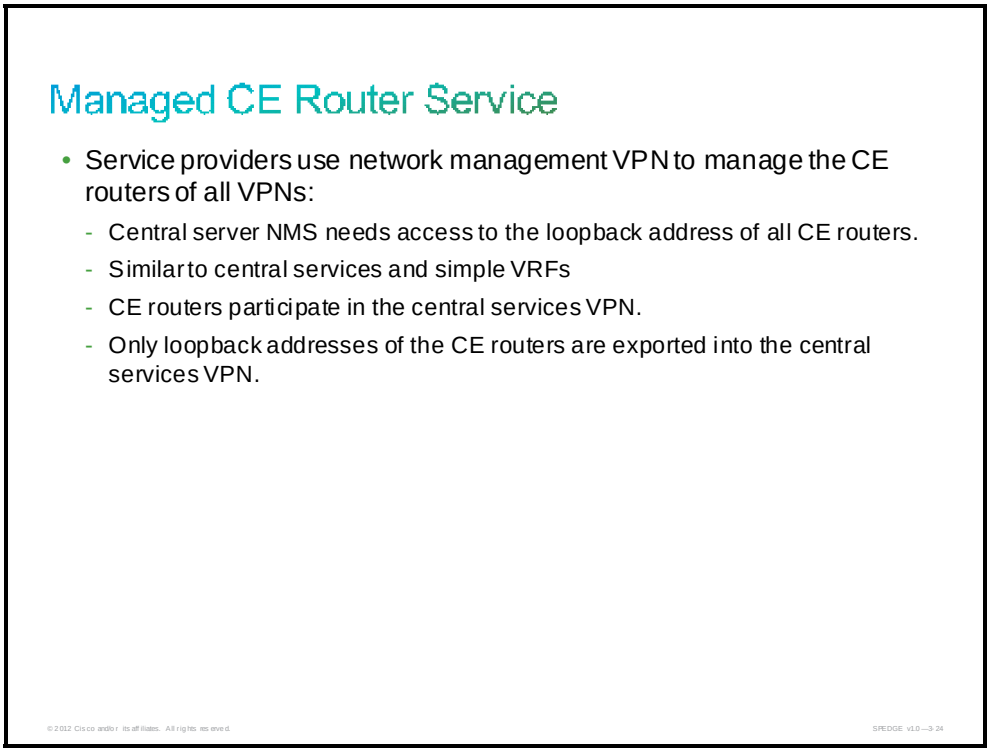
Note Depending on when you configure the **export map** command, you might need to use the **clear bgp** command to force the existing BGP session to propagate the extended communities.

In this example, routes from a certain address block are marked with an additional RT in the originating VRF and are automatically inserted into the receiving VRF based on their RT.

https://t.me/learningnets

Managed CE Router Service

This topic describes managed CE router service.



If the service provider is managing the customer routers, it is convenient to have a central point that has access to all CE routers but does not have access to the other destinations at the customer sites. This requirement is usually implemented by deploying a separate VPN for management purposes. This VPN needs to see all the loopback interfaces of all the CE routers. All CE routers need to see the network management VPN. The design is similar to that of the central services VPN; the only difference is that you mark only loopback addresses to be imported into the network management VPN.

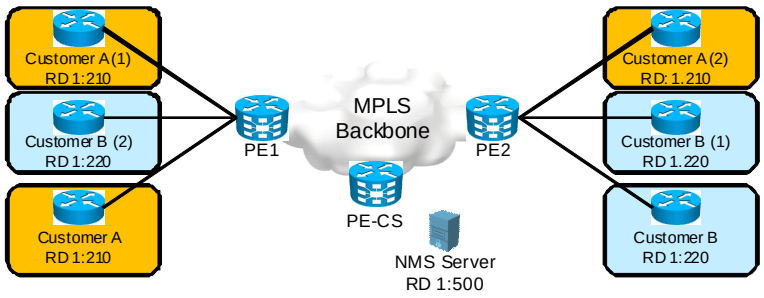
Managed CE router service features include the following:

- The central server network management system (NMS) needs access to the loopback address of all CE routers
- It is similar to central services and simple VRFs
- CE routers participate in the central cervices VPN
- Only the loopback addresses of the CE routers are exported into the central services VPN

https://t.me/learningnets

Configuration Overview

- Create one VRF per customer VPN per PE router:
  - Assign the same RD to each customer VRF.
- Create an NMS VRF on the central services PE router:
  - Assign a unique RD to the NMS VRF.

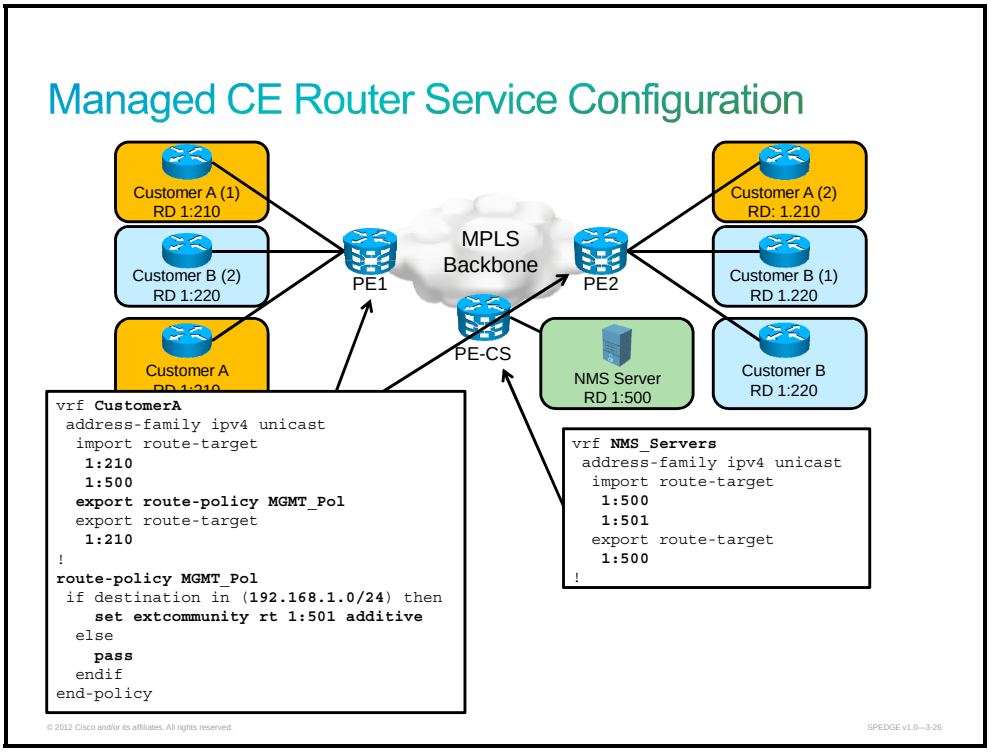


© 2012 Cisco and/or its affiliates. All rights reserved.

The configuration overview is as follows:

- Create one VRF per customer VPN per PE router
  - Assign the same RD to each customer VRF
- Create an NMS VRF on the central services PE router
  - Assign a unique RD to the NMS VRF

https://t.me/learningnets



You can have a configuration for a customer VRF with differentiated RT export for loopback addresses.

The figure shows this example. An export route policy is used to match one part of the IP address space and attach an additional RT to the routes within this address space (CE router loopback addresses).

Note	The routing protocol between PE and CE routers must be secured (with distribute lists or prefix lists) to prevent customers from announcing routes in the address space dedicated to network management; otherwise, customers can gain two-way connectivity to the network management station.

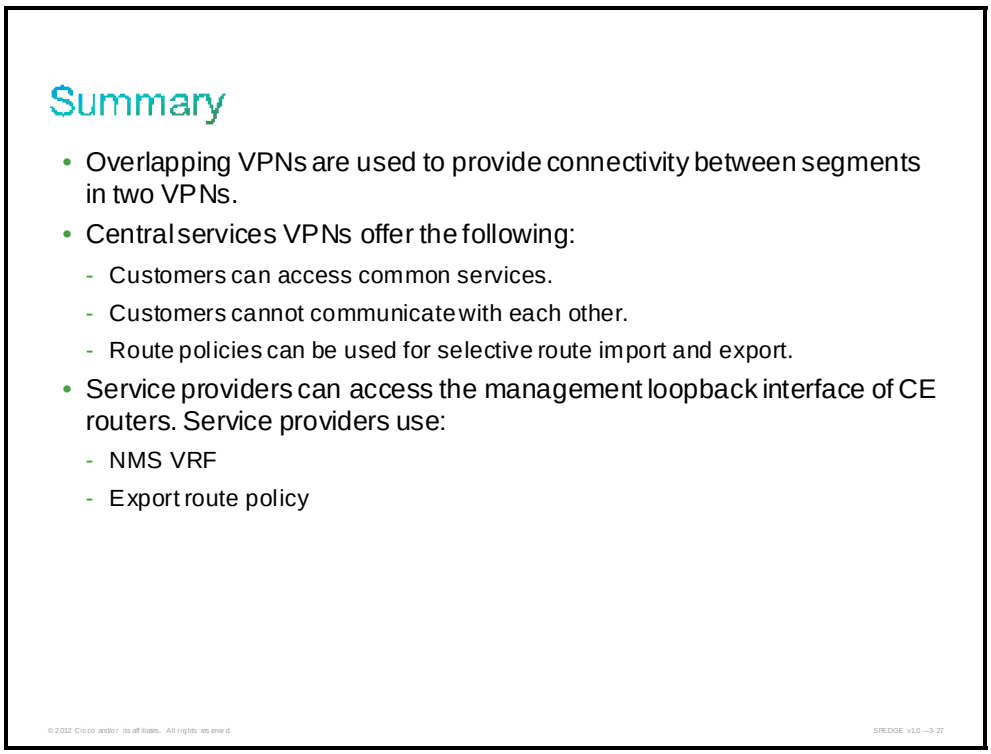
The CE router loopback addresses are then imported into the server VPN based on the additional RT attached to them during the export process.

Note	This design allows client sites to send packets to the network management VPN regardless of the source address. Special precautions should be taken to protect the network management VPN from potential threats and DoS attacks coming from customer sites.

https://t.me/learningnets

Summary

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



https://t.me/learningnets

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 2

Implementing Internet Access and MPLS Layer 3 VPNs

Overview

Integrating Internet access with a Multiprotocol Label Switching (MPLS) VPN solution is one of the most common service provider business requirements. This lesson provides a good understanding of underlying design issues, several potential design scenarios, and some sample configurations. Various topologies and implementation methods are discussed, along with ways to separate Internet access from VPN services.

Objectives

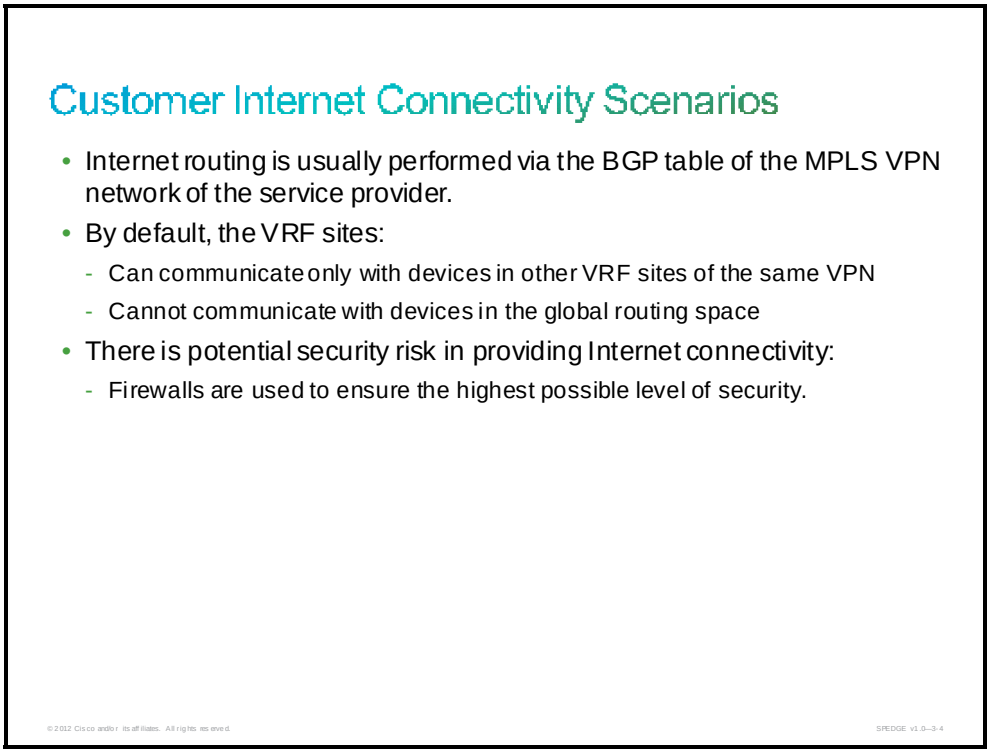
Upon completing this lesson, you will be able to describe common customer Internet connectivity scenarios and identify design models for combining Internet access with MPLS Layer 3 VPN services. You will be able to meet these objectives:

- Describe common customer Internet connectivity scenarios and identify design models for combining Internet access with MPLS Layer 3 VPN services
- Describe implementation of the Internet access service totally separate from MPLS Layer 3 VPN services
- Describe implementation of the Internet access solutions in which Internet access is provided as a separate VPN

https://t.me/learningnets

Internet Access Models with MPLS VPNs

This topic describes common customer Internet connectivity scenarios and identifies design models for combining Internet access with MPLS Layer 3 VPN services.



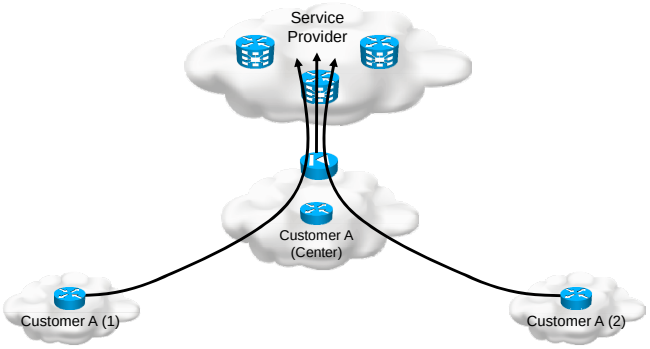
Internet routing is usually performed via the Border Gateway Protocol (BGP) table of the MPLS VPN network of the service provider. This BGP table is in the global routing space, not in the virtual routing and forwarding (VRF) context. By default, the VRF sites can communicate only with other VRF sites in the same VPN, not with anything in the global routing space. Therefore, something must be done to provide Internet access (global context) to the customer edge (CE) routers (VRF context). The following subtopics explain how to provide Internet access to VRF sites. Internet access is possible only for those customer IP subnets that are not from the private IP addressing space (RFC 1918).

Note	As soon as the VPN has Internet connectivity, a potential security risk exists. Take the proper steps—such as filtering and using a firewall—to ensure the highest possible level of security.

https://t.me/learningnets

Classical Internet Access

- Customer connects to the Internet through a central site firewall:
  - Deals with security issues
  - Provides NAT or proxy services as needed
- Internet traffic goes across the central site:
  - Traffic flow is not optimal.



© 2012 Cisco and/or its affiliates. All rights reserved.

Classical Internet access is implemented through a central firewall that connects the customer network to the Internet in a secure fashion.

The customer network and the Internet are connected only through the firewall. The addressing requirements for this type of connection are simple. The customer is assigned a small block of public address space that is used by the firewall. The customer typically uses private addresses inside the customer network if they are using IPv4 addressing. The firewall performs Network Address Translation (NAT) between the private addresses of the customer and the public addresses that are assigned to the customer by the ISP. Alternatively, the firewall might perform an application-level proxy function that also isolates private and public IP addresses.

If a customer is using IPv6 addressing, there is no need for NAT, but the customer still needs other functions that a firewall provides.

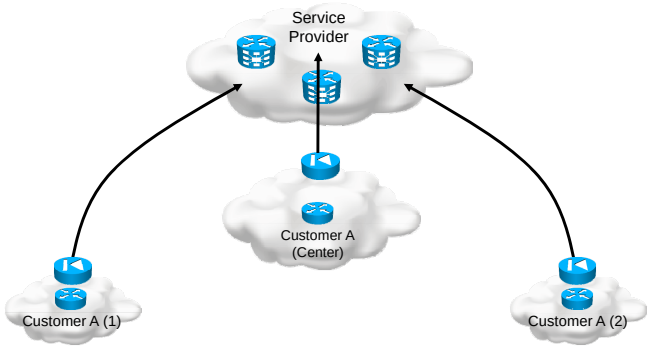
Several benefits are associated with this design. The setup is well known, and the expertise needed to implement it is simple and straightforward. Only one interconnection point between the secure customer network and the Internet needs to be managed.

The major drawback of this design is the traffic flow. All traffic from the customer network to the Internet passes through the central firewall. Although this flow might not be a drawback for smaller customers, it can be a severe limitation for large organizations with many users, especially when they are geographically separated.

https://t.me/learningnets

Multisite Internet Access

- Customers have Internet access directly from every site.
- Optimum traffic flow for Internet traffic
- Each site has to deal with security issues:
  - Managed firewall offered by service provider
  - Customer firewall



© 2012 Cisco and/or its affiliates. All rights reserved.

Some customers find the traffic flow limitations of the central firewall setup too limiting. To bypass the limitations of Internet access through a central firewall, some customers use designs in which each customer site has its own independent Internet access.

This design solves traffic flow issues, but the associated drawback is higher exposure. Each site needs to be individually secured against unauthorized Internet access, leading to the increased complexity of managing a firewall at every customer site.

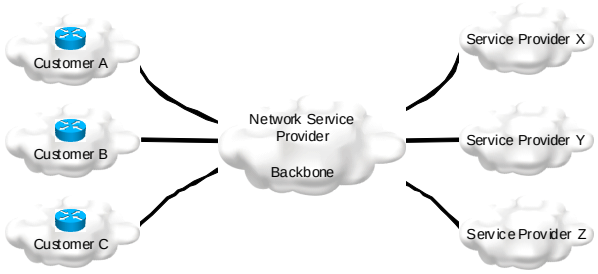
To achieve Internet access from every customer site, each CE router must forward VPN traffic toward other customer sites and forward Internet traffic toward Internet destinations. The two traffic types are usually sent over the same physical link, but over different logical links, to minimize costs.

For customers that do not want the complexity of managing their own firewalls, a managed firewall service offered by the service provider can help address the security issues of Internet connectivity.

https://t.me/learningnets

Wholesale Internet Access

- Customer chooses an ISP and selects services.
- User can access different services offered by different service providers.
- Internet access backbone:
  - Provided by NSP
  - Used to interconnect customer with service provider



© 2012 Cisco and/or its affiliates. All rights reserved.

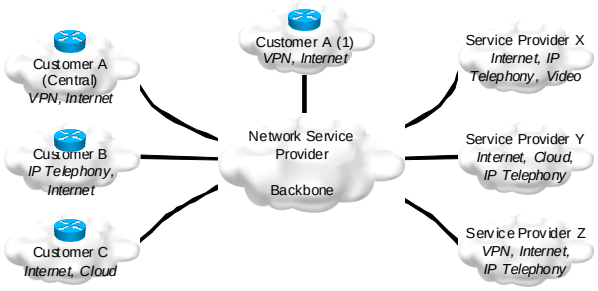
A service provider might provide wholesale Internet access from a range of upstream ISPs to satisfy the connectivity and reliability requirements of various customers.

The selection of upstream ISPs and the corresponding configuration processes should therefore be as easy as possible for the service provider. From an Internet perspective, customers A, B, and C are connected to ISP X or ISP Y. The IP address space that the customer uses should be allocated from the block of addresses that is administered by the selected ISP. The service provider that provides wholesale Internet access might need to use a different address for each upstream ISP.

https://t.me/learningnets

Service Provider Shared Backbone

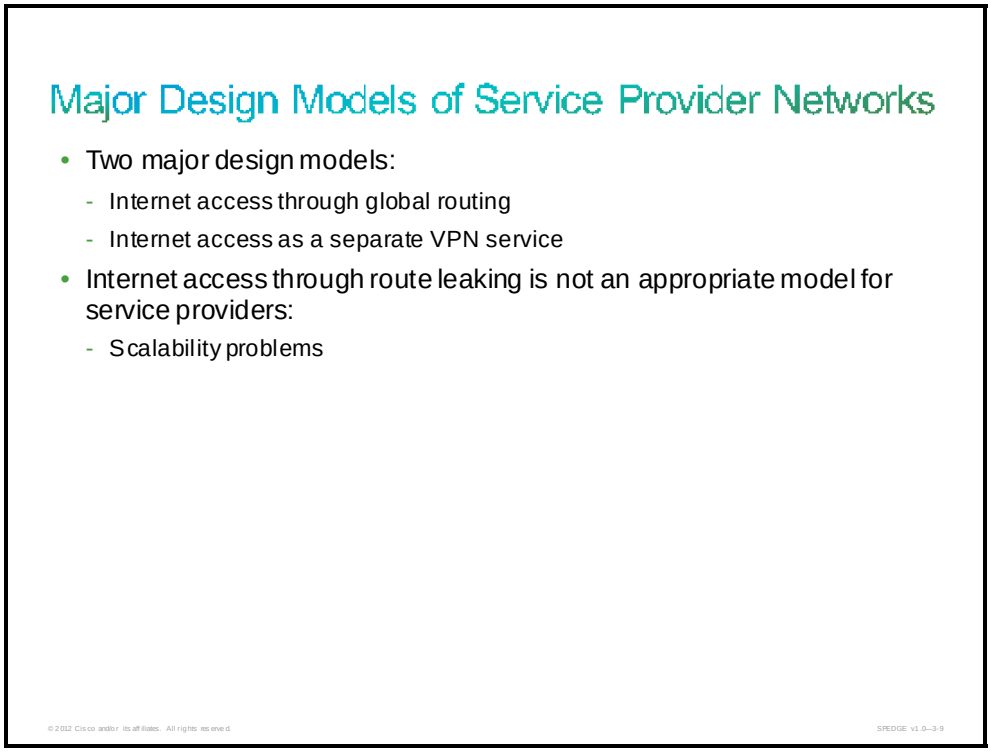
- Internet used to be the most popular service.
- Clients expect different services from their service providers now:
  - Internet, video, IP telephony, cloud services, and so on
- Cisco IP NGN architecture supports multiple services in a common backbone.



© 2012 Cisco and/or its affiliates. All rights reserved.

Because Internet access is one of the most popular services that service providers offer their customers, many service providers offer Internet access as well as MPLS VPN service on their shared backbone. Integrating Internet access with an MPLS VPN solution is one of the most common service provider business requirements. A background of common customer Internet connectivity scenarios will help in assessing possible implementations.

https://t.me/learningnets

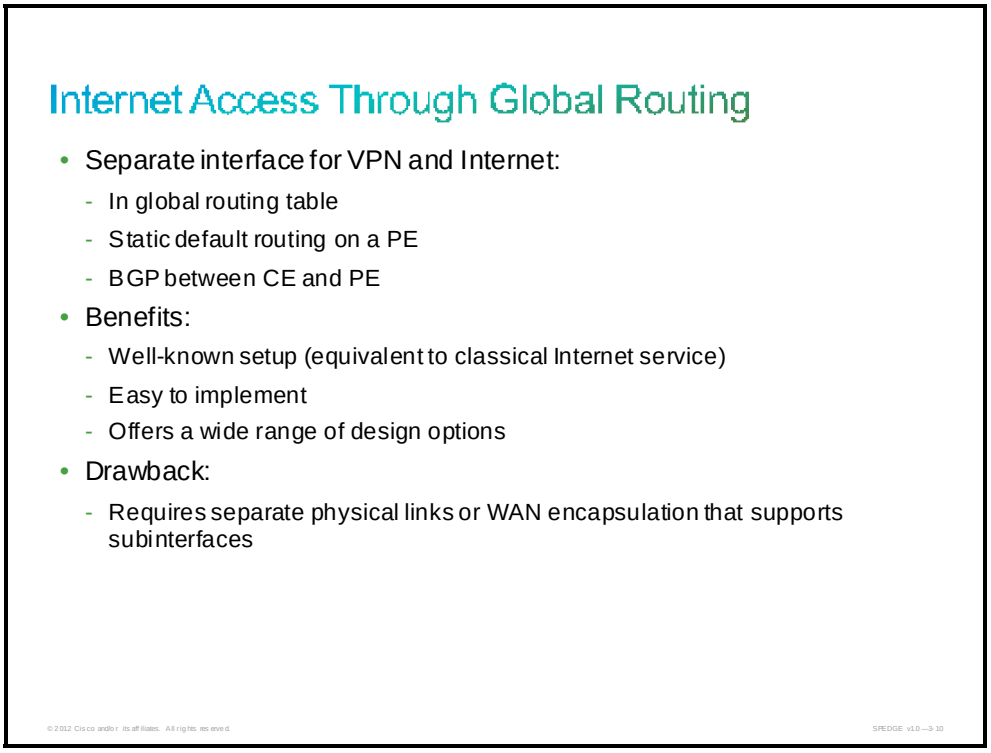


Network designers who want to offer Internet access and MPLS VPN services on the same backbone can choose between these two major design models:

- Internet access that is implemented through global routing on the provider edge (PE) routers and that is not a VPN service
- Internet access that is implemented as a separate VPN in the ISP network

In both cases, security should be the most important concern for customers when they connect to public networks. Customers should isolate private VPNs from Internet traffic, either physically (on a separate interface) or on a subinterface. Appropriate firewall support, either in a dedicated device or integrated in the router Cisco IOS Software, is a necessity. Depending on the network addressing, NAT will be needed for most customers.

https://t.me/learningnets

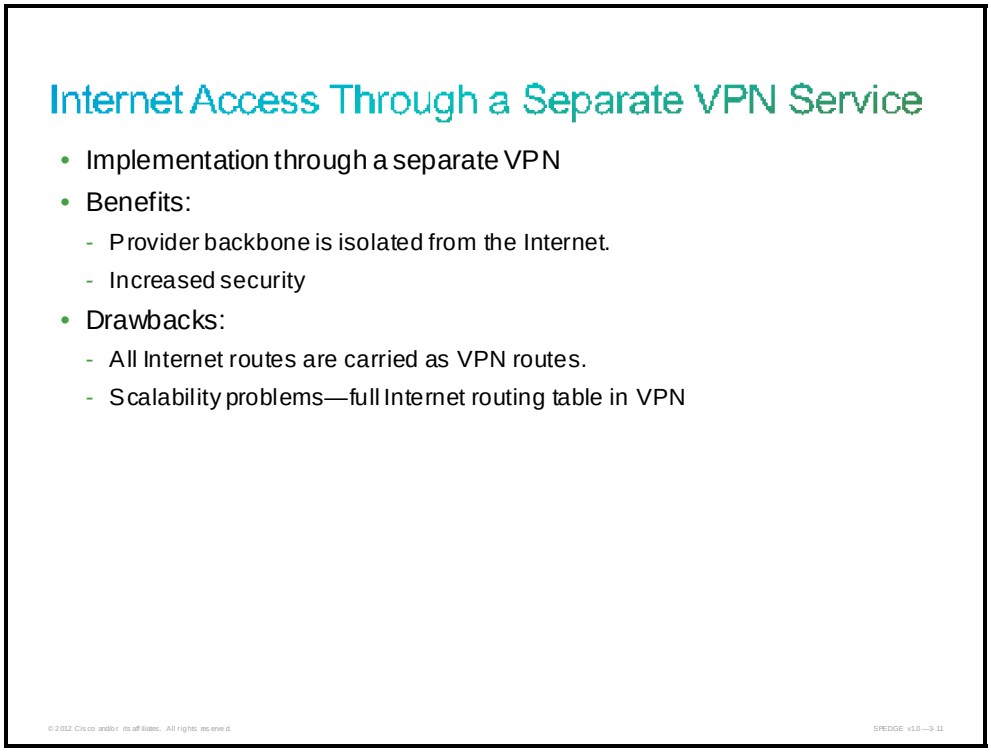


Implementing Internet access through global routing is identical to building a traditional IP backbone that offers Internet services. Depending on whether the customer sites need full Internet routes, static routes or BGP is used for routing to the Internet. For full Internet routes, BGP is deployed between the PE routers and the PE Internet gateway to exchange Internet routes, and the global routing table on the PE routers is used to forward the customer traffic toward Internet destinations. The PE routers might or might not have the full Internet routing table. The PEs and the provider Internet gateway are in the same Internal BGP (IBGP) area.

The PE routers also use Multiprotocol BGP (MP-BGP) to support VPNs for their customers. The customers reach the global routing table by using a separate logical link for Internet access. Internet access through global routing on separate logical links is easy to set up; it is the equivalent of the classical combination of Internet and VPN services that many customers use today. This setup is also compatible with all the Internet services required by some customers (for example, the requirement to receive full Internet routing from a service provider).

The drawback of this design is the increased complexity, or cost, of the PE-CE connectivity. Separation of Internet and VPN connectivity requires either two separate physical links or a single physical link with encapsulation that supports subinterfaces; for example, Frame Relay, PPP or using Ethernet links and 802.1Q encapsulation.

https://t.me/learningnets

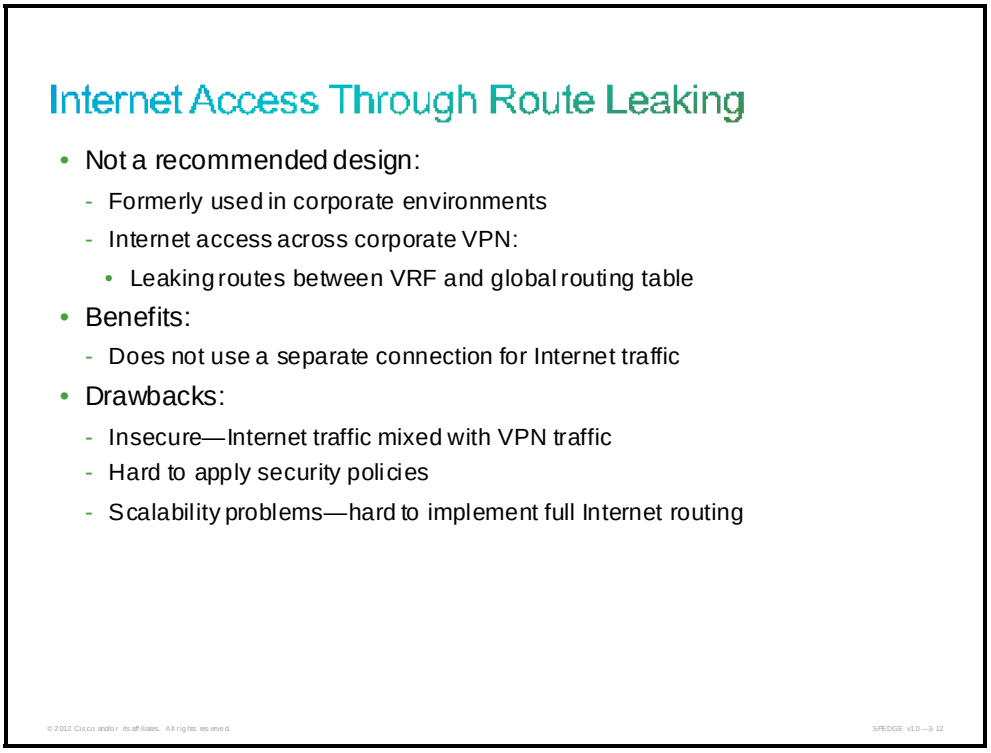


For a service provider, implementing Internet access through a separate VPN is similar to offering another managed VPN service.

The major benefit of implementing Internet access as a separate VPN is increased isolation between the provider backbone and the Internet—which results in increased security for the provider. The flexibility of MPLS VPN topologies also provides for some innovative design options that allow service providers to offer services that were simply impossible to implement with pure IP routing.

The obvious drawback of running the Internet as a VPN in the MPLS VPN architecture is the scalability of such a solution. An Internet VPN cannot carry full Internet routing because of the scalability problems that are associated with having all the Internet routes inside a single VPN.

https://t.me/learningnets



Another variant for Internet support with MPLS VPNs is to provision route leaking into the global routing context. Although not a recommended design, this option is briefly discussed to show an alternate practice that has been used in the industry.

Some customers might want to obtain Internet access across their corporate VPN by leaking routes between the VRF and global routing tables.

Caution	For security reasons, this approach is not recommended. Bringing in Internet traffic by using the corporate VPN is not a good practice and negates the isolation of the corporate VPN.

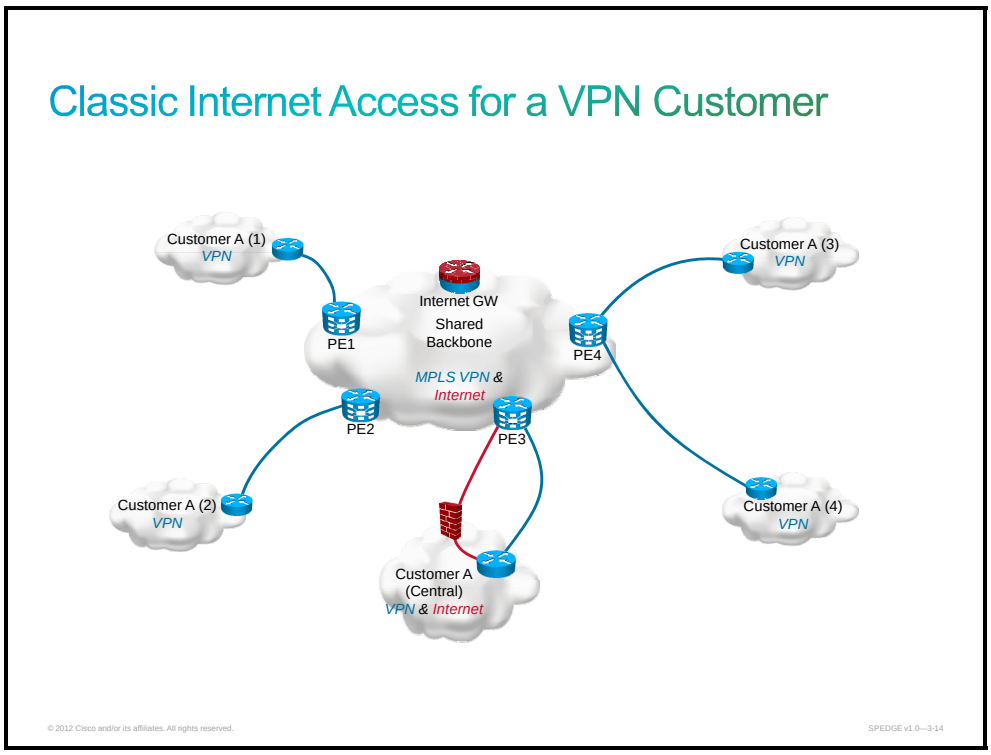
With route leaking, the customer site uses a static default route in the VRF table that points to the global next-hop address of an Internet gateway. Any packets that use the default route leave the VPN space and are routed according to the global routing table at the PE that is the next-hop router. This feature allows leaking of VPN packets into the global address space.

Note	This approach is not recommended and is not discussed further in this course. This option is discussed briefly only to show an alternate practice that has been used in the industry.

https://t.me/learningnets

Separate Internet Access and VPN Services

This topic describes implementation of the Internet access service totally separate from MPLS Layer 3 VPN services.



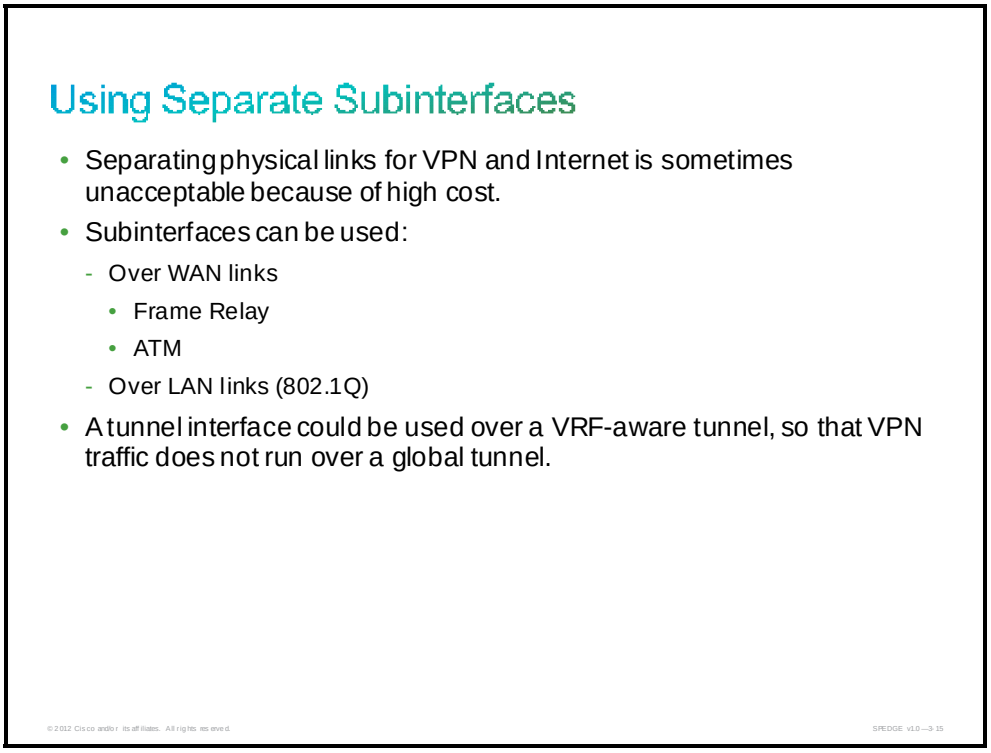
The classical Internet access design for a customer is based on a separate Internet access model. One central customer site has connectivity to the Internet and provides access to the rest of the customer sites. The central site either connects through a firewall or runs the Cisco IOS Firewall Feature Set.

In the shared service provider backbone, the PE routers have full or partial Internet routes to offer the customer. The Internet gateway of the provider is in the same IBGP domain as the provider and PE routers.

This design model can easily map to a customer with an MPLS VPN implementation. In this example, the customer network has been interconnected with an MPLS VPN. A central CE router has Internet connectivity and provides Internet access, through a firewall, for all the sites in the customer network.

This traditional Internet access implementation model provides maximum design flexibility because the Internet access is completely separated from the MPLS VPN services. However, the limitations of traditional IP routing prevent this implementation method from being used for innovative Internet access solutions, such as wholesale Internet access.

https://t.me/learningnets



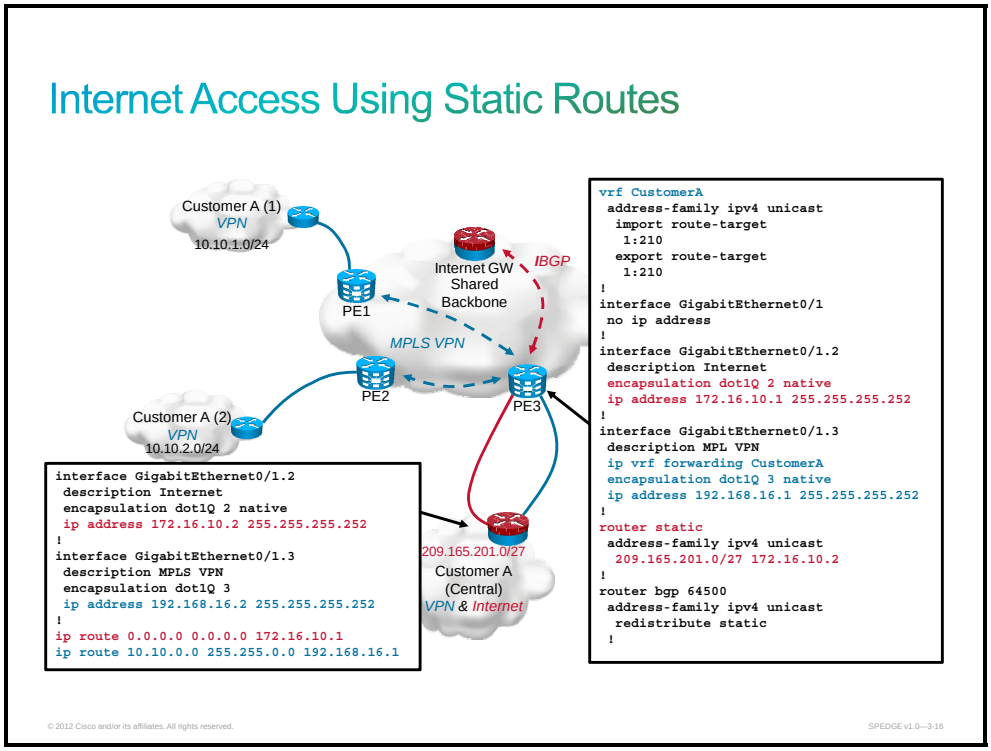
Instead of separate physical links for VPN and Internet traffic, subinterfaces can be used to create two logical links over a single physical link. Subinterfaces can be configured only on WAN links that use Frame Relay or ATM encapsulation (including xDSL) and on LAN links that use any VLAN encapsulation (802.1Q).

For other encapsulation types, a tunnel interface can be used between the CE router and the PE router. Depending on the router platform and Cisco IOS Software version, virtual routing and forwarding (VRF)-aware tunnels are now supported.

- VRF-aware tunnels remove the need for the endpoints of the tunnel to be in a global routing table.
- Without a VRF-aware tunnel, MPLS VPN traffic would need to be tunneled across the Internet interface.

Note	Further information on VRF-aware tunnels is outside the scope of this course.

https://t.me/learningnets



Using static routing on the CE and PE routers is the simplest and most common implementation for providing Internet access. The figure illustrates a configuration that is used to implement Internet access through two 802.1Q subinterfaces with static routes.

In this simple example, the Customer A router does not need to receive full Internet routing. To reach the Internet, the Customer A router just needs a default static route to PE3. PE3 has a route to the Internet through the Internet gateway, as well as a static route for the customer A subnets that point to the Customer A router. The full Internet routing table needs to be present only on the Internet gateway.

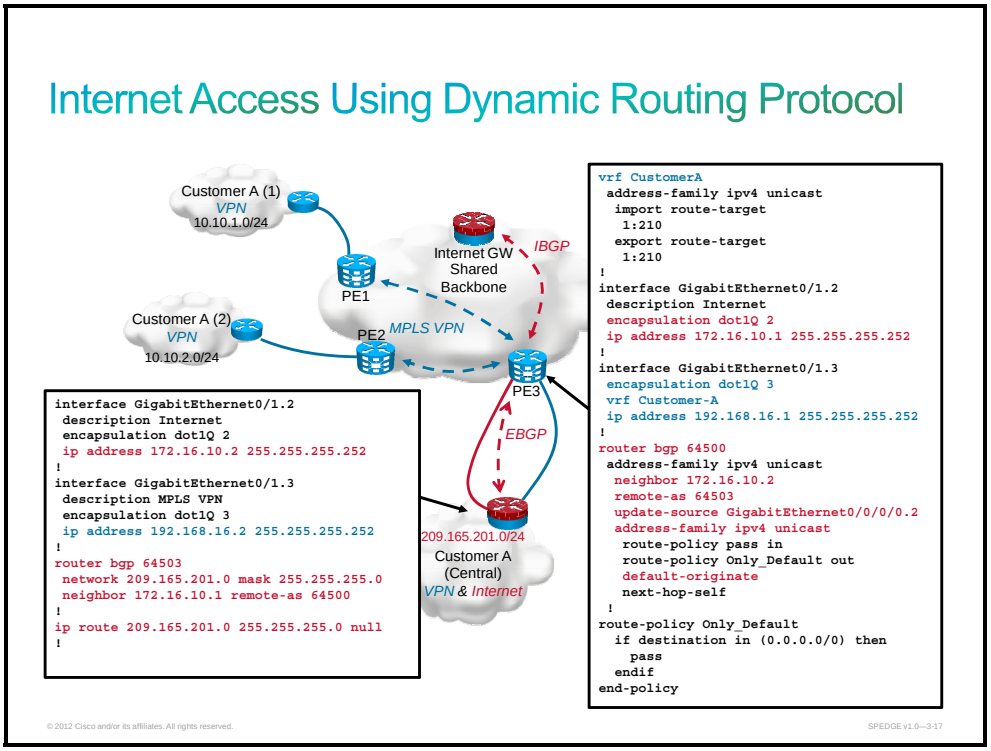
The following configuration steps are performed:

1. The customer VRF instance is created for private MPLS VPN.
2. The VPN subinterface is created and associated with the proper VLAN. The subinterface is added to the customer VRF.
3. Another subinterface is created for Internet access.
4. On the PE router, configure static routes for customer public address space pointing to customer next hop. Redistribute these routes in BGP routing protocol.

In addition to the PE configuration, the CE router implements a default static route that points to the PE router.

Note	On the CE-Central router, distribution of the default route might be needed so that remote sites can also access the Internet. Issues of security and private addresses would need to be resolved.

https://t.me/learningnets



In this example, Customer A is using a dynamic routing protocol to establish Internet connectivity.

BGP routing protocol is configured on the Customer A router. Customer A advertises its own public address space to External BGP (EBGP) neighbor. It is important that customer public address space is inserted into the customer routing table. If it is not in the routing table, the network is not advertised using BGP.

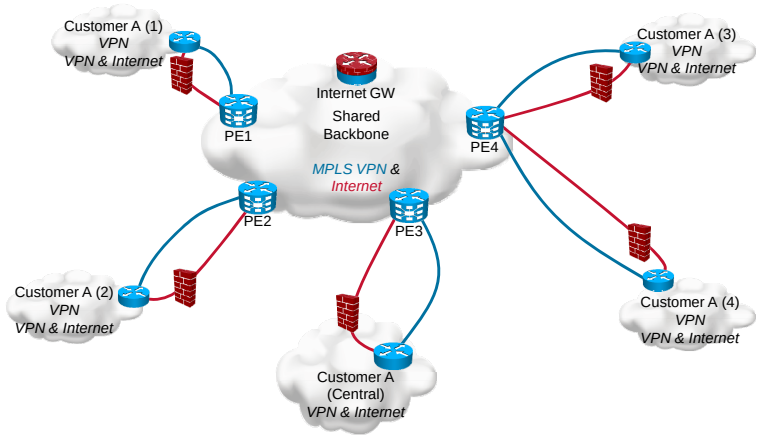
On the PE router, the service provider has to configure the VRF instance and subinterface for the private MPLS VPN of the customer and another subinterface for Internet access.

BGP is used for route exchange between the customer and service provider. The service provider can originate the default route to the customer peer. The service provider can also filter advertised routes and send only the default route to the customer. A route policy is used to filter routes on Cisco IOS XR routers, and prefix lists are used to filter advertised routes on IOS routers.

https://t.me/learningnets

Multisite Internet Access

- Every CE router needs two links (or subinterfaces).
- Complex network setup
- Expensive solution

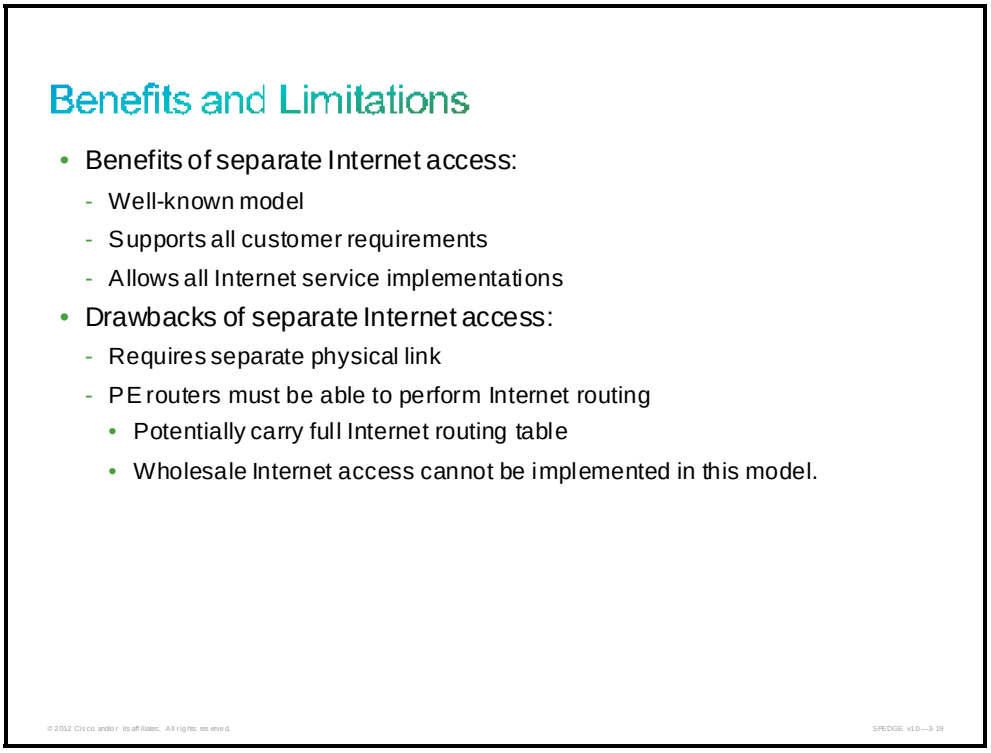


© 2012 Cisco and/or its affiliates. All rights reserved.

Another option is to provide separate Internet access at every customer site. In this case, two physical (or logical) links between every CE router and its PE router would be needed. This design often becomes too complex or too expensive to implement. Issues such as customer route propagation to the Internet and securing access at multiple access points would need to be resolved.

Note	The allowas-in feature might need to be configured on the PE router if the customer is propagating individual site routes to the Internet through BGP.

https://t.me/learningnets



The benefits of a separate Internet access design model are as follows:

- The model is well known and widely understood.
- The model supports all customer requirements, including multihomed customer connectivity with full Internet routing.
- The model allows all Internet service implementations, including BGP sessions with customers.

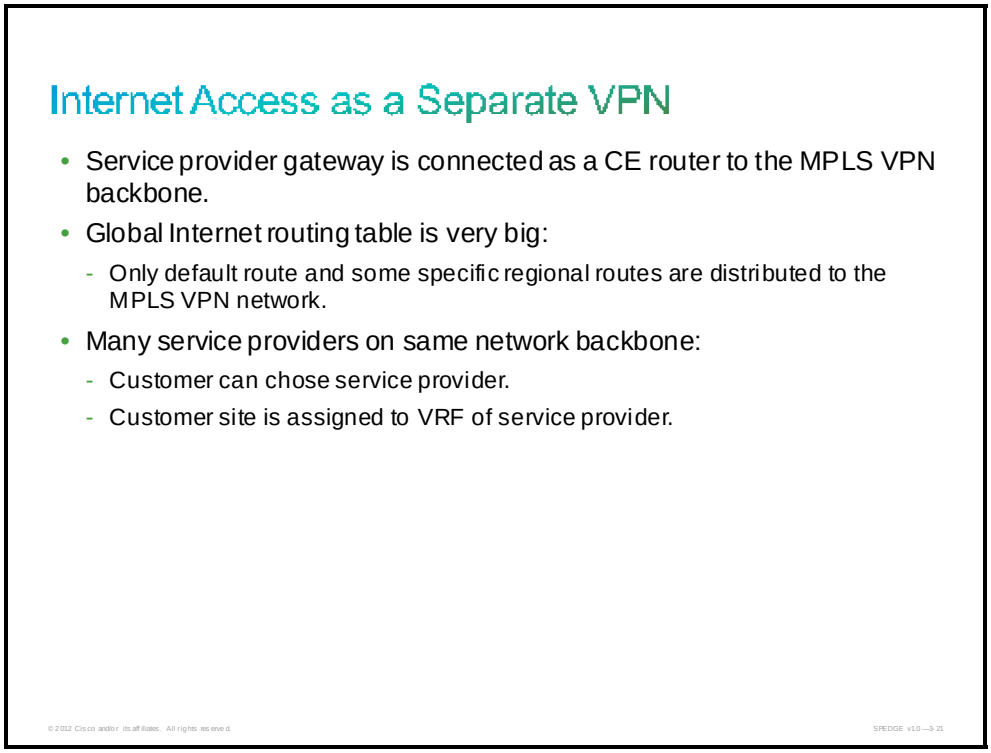
The drawbacks of this model are as follows:

- The model requires two dedicated physical links between the PE and the CE router or specific WAN or LAN encapsulations that might not be suitable for all customers.
- The PE routers must be able to perform hop-by-hop Internet routing and either use the default route to reach the Internet or carry the full Internet routing table.
- Advanced Internet access services (centrally managed firewall service or wholesale Internet access service) cannot be realized with this model.

https://t.me/learningnets

Internet Access as a Separate VPN

This topic describes implementation of the Internet access solutions in which the Internet access is provided as a separate VPN.



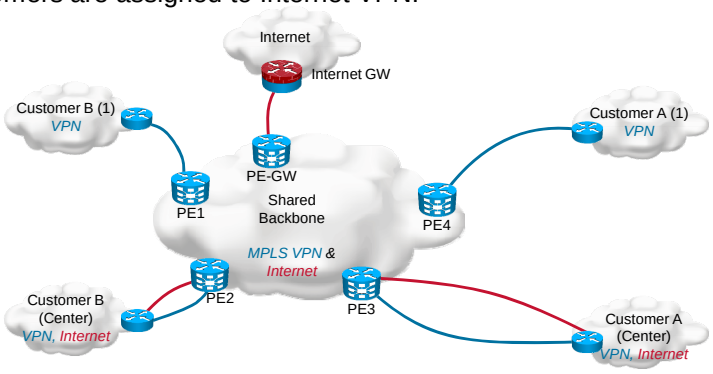
The MPLS VPN architecture can provision a separate VPN to provide Internet access for VPN customers. The service provider defines the Internet VPN and can use different MPLS VPN topologies to implement various types of Internet access. Under this design model, the provider Internet gateways appear as CE routers to the MPLS VPN backbone. Customer Internet access is enabled by using a dual VPN topology that supports both an Internet VPN and a customer VPN across separate customer interfaces.

In this design, the Internet VPN should not contain the full set of global Internet routes because that would make the solution completely nonscalable. The provider Internet gateway routers should announce a default route toward the PE routers. To optimize local routing, the local and regional Internet routes should be inserted in the Internet VPN.

https://t.me/learningnets

Internet Gateway Configuration

- Internet gateway has full Internet routing table:
  - Only subset of all routes sent to customers
- Internet gateway acts as a CE router.
- Internet VPN is used for Internet access.
- Customers are assigned to Internet VPN.

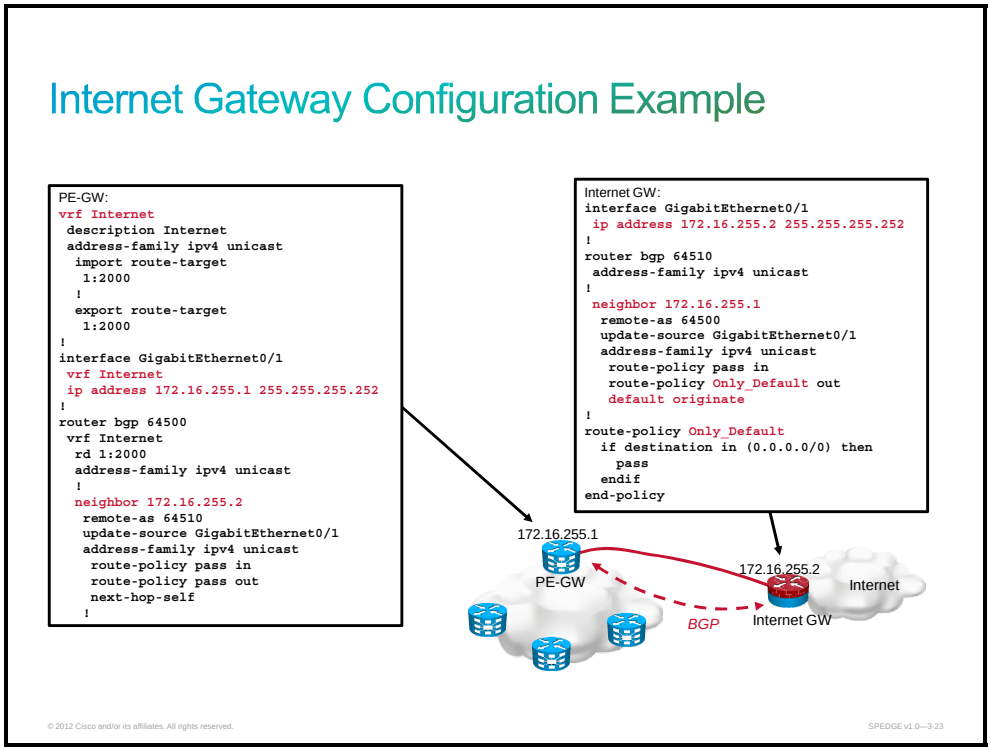


© 2012 Cisco and/or its affiliates. All rights reserved.

When the service provider implements Internet access as a separate VPN, the Internet backbone is carried on a VPN, which is isolated from the provider backbone. This topology results in increased security for the provider backbone because Internet hosts can reach only PE routers, not the core provider routers. The VPN customers are connected to the Internet simply through an additional VRF instance at the PE.

Internet gateway acts as a CE router and holds the full BGP routing table. Only the subset of routes and the default route are advertised to clients assigned to Internet VPN.

https://t.me/learningnets



The figure shows a sample configuration of the Internet gateway router with default route advertisement.

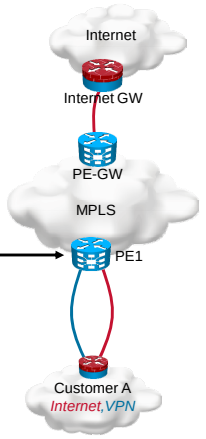
The Internet gateway should be able to route traffic to the Internet. An EGBP session is established with the PE gateway router. Only the default route is advertised to the PE gateway router.

On the PE gateway router, a new VRF instance for Internet is used. The interface facing the Internet gateway is assigned to the Internet VRF. In the BGP process, the Internet VRF has to be enabled and appropriate address families have to be activated (such as IPv4 and IPv6). The PE gateway router distributes the default Internet route among the other PE routers in the MPLS VPN network.

https://t.me/learningnets

Customer Internet Access from Central Site

```
PE1:
vrf CustomerA
 address-family ipv4 unicast
  import route-target
   1:210
  export route-target
   1:210
!
interface GigabitEthernet0/1.2
 description Internet
 encapsulation dot1Q 2
 ip address 172.16.10.1 255.255.255.252
!
router bgp 64500
 address-family ipv4 unicast
 !
 address-family vpnv4 unicast
 !
!
 vrf Internet
  rd 1:2000
  neighbor 172.16.10.2
  remote-as 64503
  address-family ipv4 unicast
   network 0.0.0.0/0
!
```



© 2012 Cisco and/or its affiliates. All rights reserved.

The classical Internet access model can easily be implemented with the Internet VPN over the MPLS VPN backbone. The link between a PE router and the Internet gateway router is assigned to the Internet VRF, as discussed previously. The Internet gateway announces a default route to the Internet.

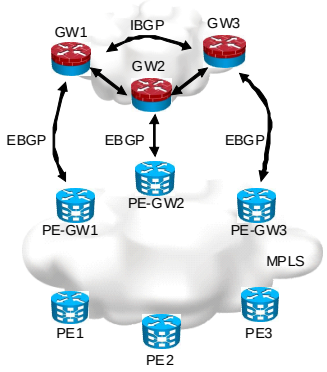
One link between the PE router and each central customer router is assigned to the customer VRF, and one is assigned to the Internet VRF.

In this example, PE1 connects the Customer A router to both the Customer A VRF and the Internet VRF.

https://t.me/learningnets

Redundant Internet Access

- All Internet gateways advertise routes.
- Internet gateways are connected to the same VRF.
- BGP metric is used to select the best route to the Internet.
- MED is used to define the primary Internet gateway.



© 2012 Cisco and/or its affiliates. All rights reserved.

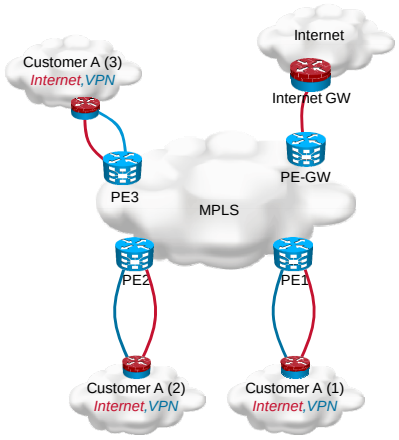
Redundant Internet access is easy to achieve when the Internet service is implemented as a VPN in the MPLS VPN backbone:

- Multiple Internet gateways (acting as CE routers) need to be connected to the MPLS VPN backbone to ensure router and link redundancy.
- All Internet gateways advertise the default route to the PE routers, resulting in routing redundancy.
- The Internet gateways also announce local Internet routes. Because these routes are announced with different BGP attributes—most notably multi-exit discriminator (MED)—the PE routers select the proper Internet gateway router as the exit point toward those destinations.
- The MED attribute can also be used to indicate the preferred default route to the PE routers. In this setup, one Internet gateway router acts as a primary Internet gateway, and the other Internet gateway router acts as a backup.
- The redundancy that has been established so far covers the path between customer sites and the Internet gateway routers. A failure in the Internet backbone might break the Internet connectivity for the customers if the Internet gateway routers announce the default route unconditionally. Conditional advertisement of the default route is therefore configured on the Internet gateway routers, which announce the default route to the PE routers only if the Internet gateway routers can reach an upstream destination.

https://t.me/learningnets

Multisite Customer Internet Access

- Internet VRF is configured on every location.
- Adds complexity
- Firewall on every site:
  - Managed firewall can be used.



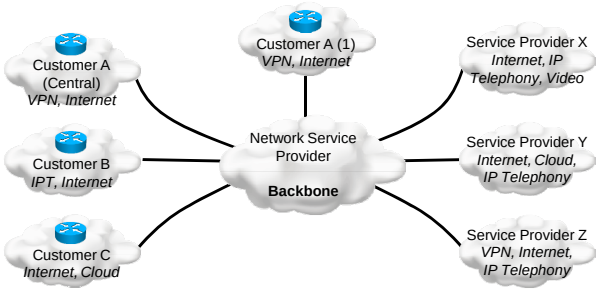
© 2012 Cisco and/or its affiliates. All rights reserved.

Multisite customer Internet access can be implemented by configuring the Internet VRF at every location. This solution adds complexity for the customer because firewall and Network Address Translation (NAT) support might be needed at every site, unless the service provider offers a central managed firewall service.

https://t.me/learningnets

Wholesale Internet Access

- A separate VPN is created for each upstream ISP.
- Each ISP gateway announces the default route to the VPN.
- Customers are assigned into the right VRF:
  - VRF assignment corresponds to ISP selection.
- ISP change is easy for administrator:
  - Only VRF has to be changed.

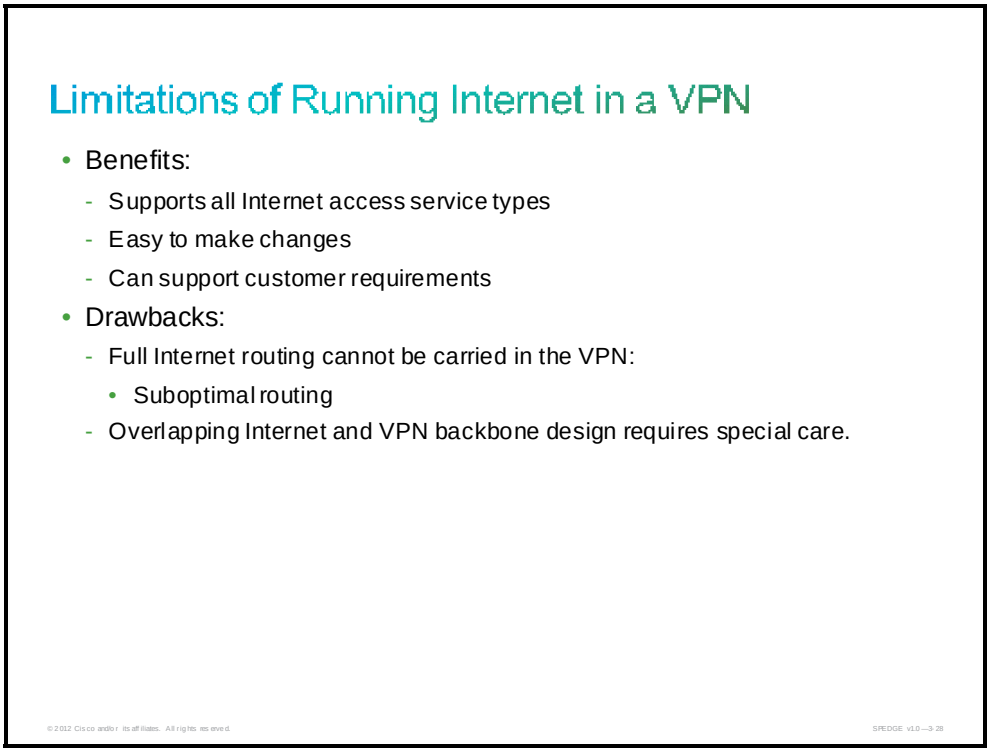


© 2012 Cisco and/or its affiliates. All rights reserved.

Wholesale Internet access is implemented by creating a separate VPN for every upstream ISP. Acting as a CE router toward the MPLS VPN-based Internet access backbone, the Internet gateway of the upstream ISP announces a default route, which is used for routing inside the VPN.

Customers are tied to upstream service providers simply by placing the PE-CE link into the VRF that is associated with the upstream service provider. Changing an ISP becomes as easy as reassigning the interface into a different VRF and managing address allocation issues.

https://t.me/learningnets



Internet access that is implemented as a separate VPN has the following benefits:

- This design model supports all Internet access services, ranging from traditional Internet access to innovative services such as wholesale Internet access.
- This design also supports all customer requirements, including full Internet routing on customer routers through an EBGP multihop session with the Internet gateway.

Internet access that is implemented as a separate VPN has the following drawbacks:

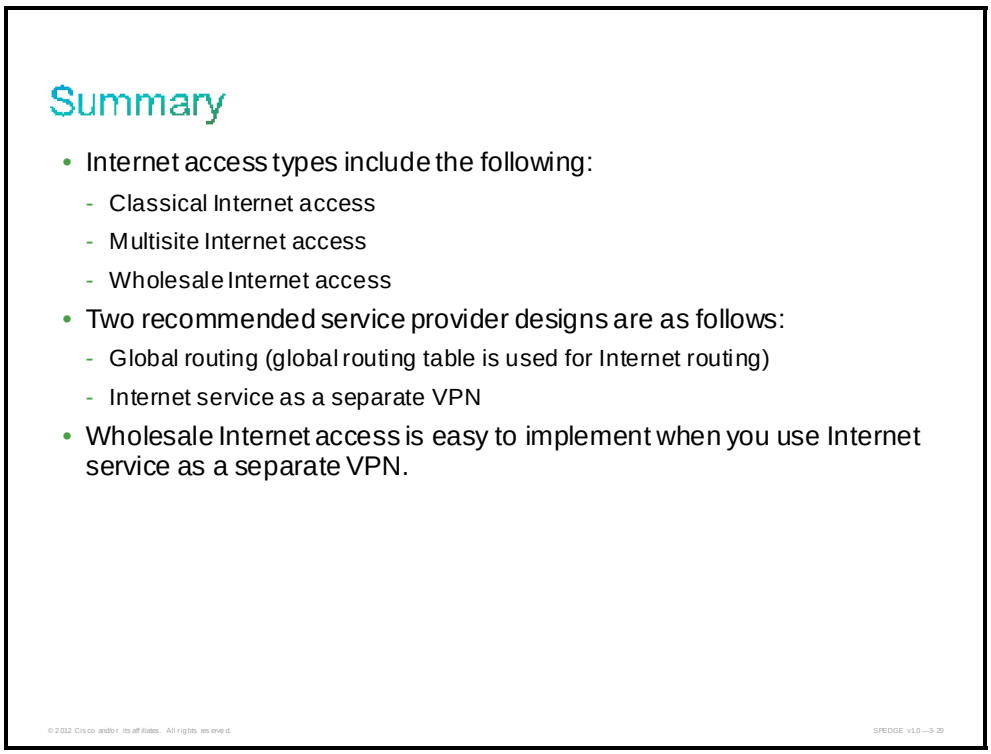
- Full Internet routing cannot be carried inside a VPN; therefore, default routing toward the Internet gateways needs to be used, potentially resulting in suboptimal routing.
- The Internet backbone gateway router is positioned as a CE router connected to the MPLS VPN backbone. If the service provider runs Internet service and MPLS VPN service on the same set of routers, the interconnection between the two services requires special considerations.

The benefits of the separate VPN design far outweigh the limitations.

https://t.me/learningnets

Summary

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



https://t.me/learningnets

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 3

Introducing MPLS Interdomain Solutions

Overview

Deployments of Multiprotocol Label Switching (MPLS) have become routine in large-scale global networks, which demand solutions to complex business and network problems. There are two primary components of the Cisco IOS MPLS Interdomain Solution: the interautonomous system (inter-AS) and Carrier Supporting Carrier (CSC).

Objectives

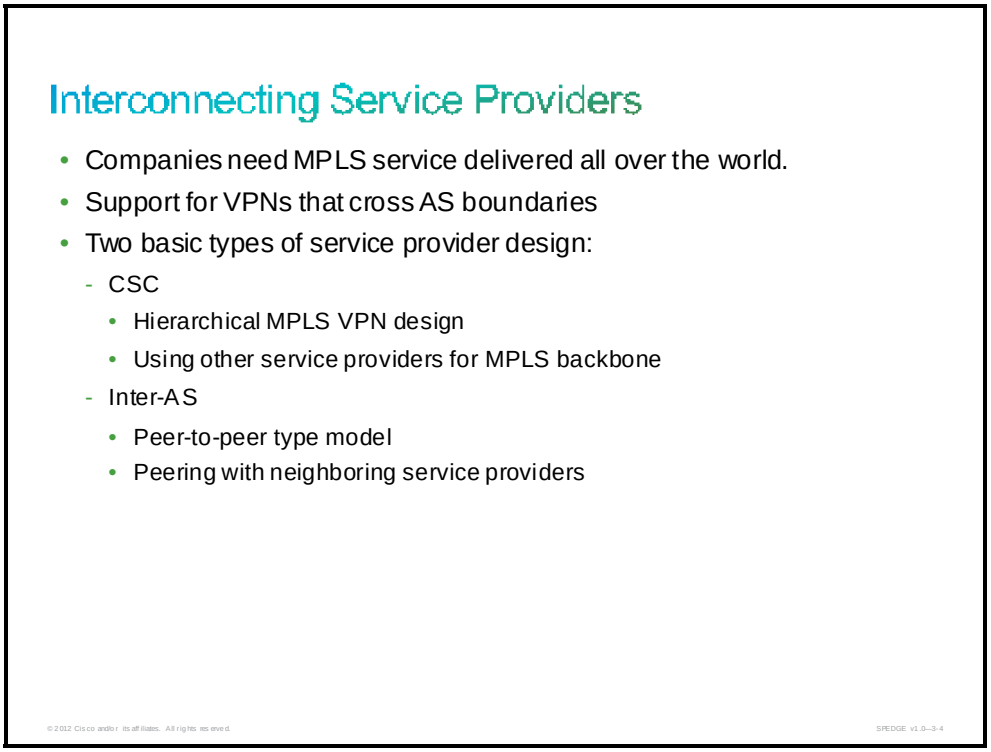
Upon completing this lesson, you will be able to introduce the Cisco IOS MPLS Interdomain solutions with different implementation options. You will be able to meet these objectives:

- Describe MPLS interdomain solutions
- Describe the CSC feature
- Describe inter-AS MPLS models

https://t.me/learningnets

MPLS Interdomain Solutions

This topic describes MPLS interdomain solutions.



Deployments of MPLS have become routine in large-scale global networks, which demand solutions to complex business and network problems. There are two primary components of the Cisco IOS MPLS Interdomain Solution: inter-AS and CSC.

Inter-AS is a peer-to-peer type model that allows the extension of VPNs through multiple provider or multidomain networks. This solution enables service providers to peer up with one another and offer end-to-end VPN connectivity over extended geographical locations for those subscribers who may be out of reach for a single provider.

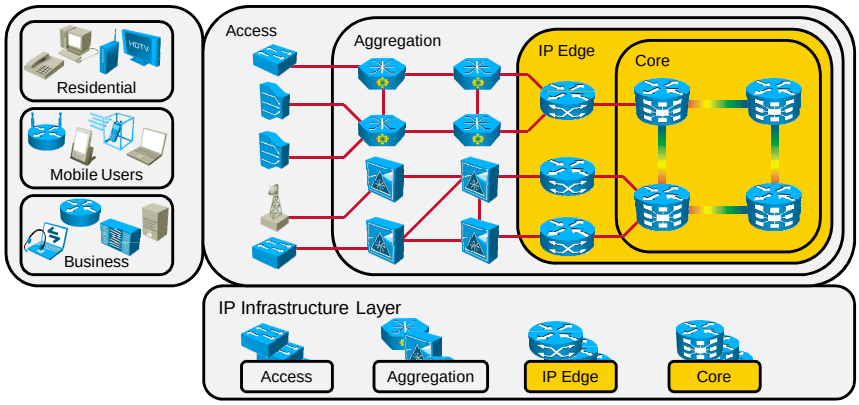
CSC is a hierarchical VPN model that allows small service providers, or customer carriers, to interconnect their IP or MPLS networks over an MPLS backbone. This eliminates the need for customer carriers to build and maintain their own MPLS backbone.

Both inter-AS and CSC can construct scalable networks that help maintain network segmentation based on internal organizational or operational boundaries.

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Carrier Supporting Carrier in Cisco IP NGN Infrastructure Layer



- MPLS interdomain solutions are part of the Cisco IP NGN infrastructure layer.
- IP edge devices run MPLS, BGP, or IGP.
- IP core devices run MPLS.

© 2012 Cisco and/or its affiliates. All rights reserved.

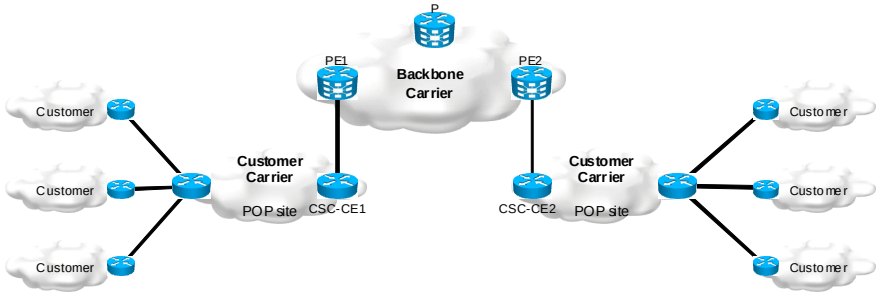
CSC is configured in the service provider core and edge network. It is part of the Cisco IP Next-Generation Network (NGN) infrastructure layer.

IP core devices run MPLS and IP edge devices run MPLS, Border Gateway Protocol (BGP), and some interior gateway protocol (IGP) routing protocols.

https://t.me/learningnets

CSC Overview

- Hierarchical MPLS VPN:
  - Backbone provider–first-level service provider
  - Customer carrier–second-level service provider
- CSC provides MPLS VPN service to other service providers.
- A large service provider acts as the backbone for smaller service providers.
- The customer carrier can be an ISP or MPLS VPN provider.

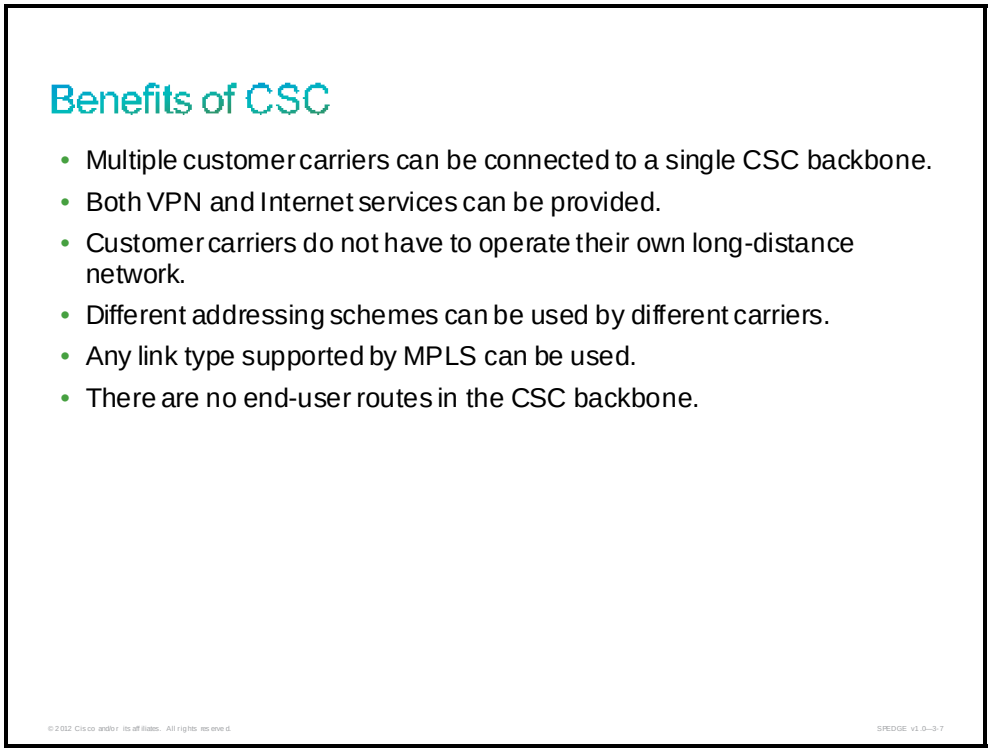


© 2012 Cisco and/or its affiliates. All rights reserved.

A carrier network carries traffic between customer sites. Large service providers can interconnect carrier networks of smaller service providers. In this scenario, a smaller service provider acts as a customer for a larger service provider.

CSC provides MPLS VPN service to other service providers. CSC creates a hierarchical structure with a first-level service provider as the backbone carrier and a second-level service provider as customer carrier. Many customer carrier sites, also called point of presence (POP) sites, can be interconnected using the backbone carrier.

https://t.me/learningnets



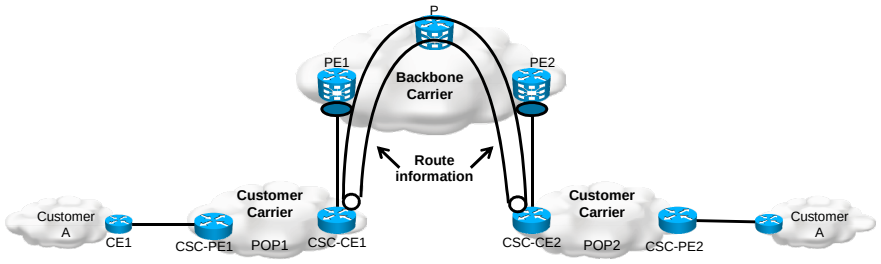
CSC has many benefits:

- A single CSC backbone can be used to connect many POP sites.
- An MPLS VPN can be used to separate traffic from the POP sites of different carriers.
- A single CSC backbone can provide services to both VPN service providers and to the ISP.
- The customer service providers do not have to operate their own long-distance network. They purchase that service from the CSC backbone carrier.
- Different customer carriers can use different addressing schemes. The customer carriers will be in separate MPLS VPNs inside the CSC backbone.
- Any link type that is supported by MPLS can be used inside the CSC backbone and as access links between the CSC backbone and customer carriers.

https://t.me/learningnets

CSC Operation

- Packets from POP1 to POP2 are propagated along a label-switched path from CE1 to CE2.
- PE and CSC-CE routers must exchange route or label information.
- Backbone carrier does not carry routing information of end customers.



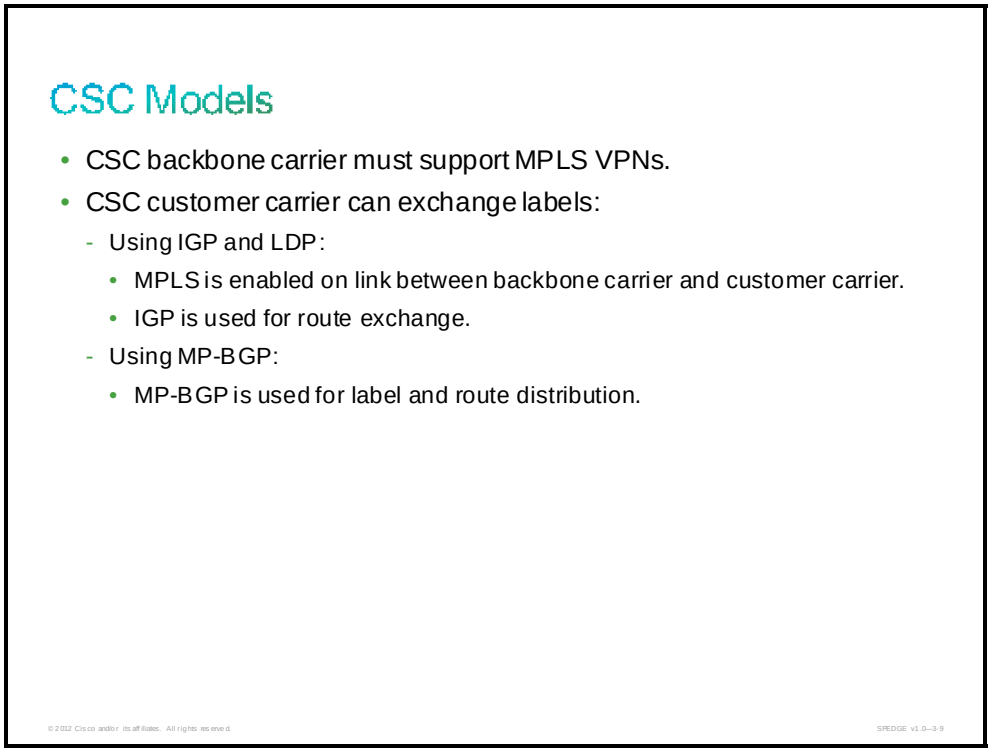
© 2012 Cisco and/or its affiliates. All rights reserved.

The CSC architecture relies on the presence of an MPLS VPN. The CSC backbone is providing an MPLS VPN service to which the customer carriers are connected as VPN sites. MPLS is used between the CSC backbone provider edge (PE) routers and the VPN sites of the customer carriers.

Virtual routing and forwarding (VRF) tables are enabled on the CSC PE routers. The label exchange between PE1 and PE2 establishes a label-switched path (LSP) from CE1 via the CSC backbone to CE2. Another LSP is also established in the other direction.

The Customer carrier can now tunnel packets between POP1 and POP2 using the LSPs. The CSC backbone does not have to know about end-user sites and their IP addresses.

https://t.me/learningnets

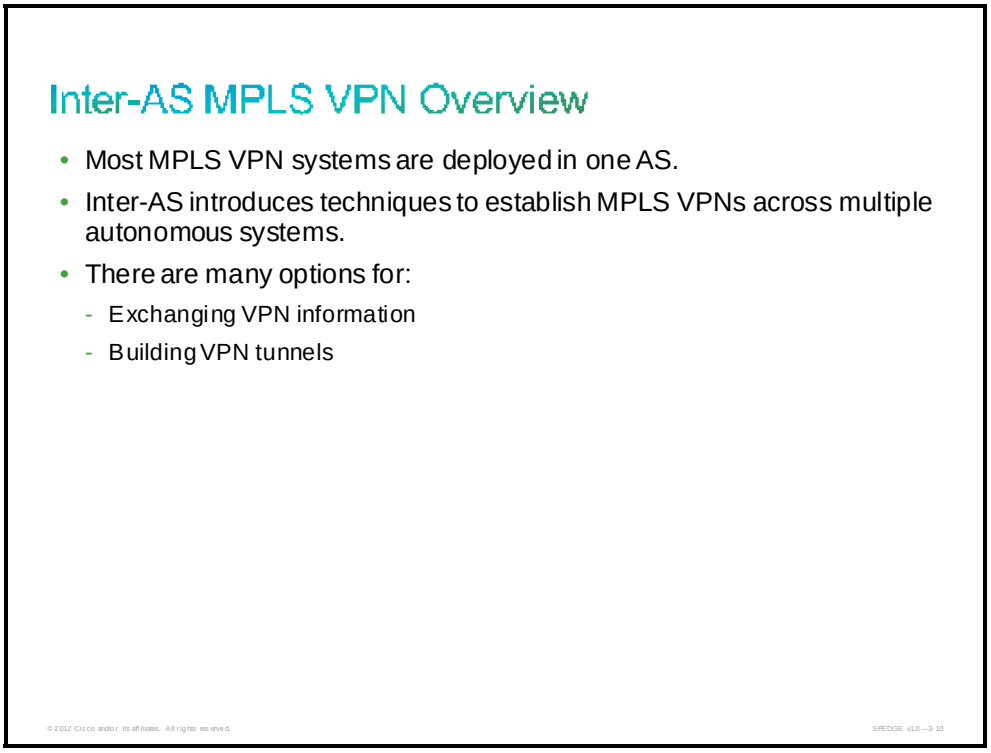


To support CSC, the CSC backbone must support MPLS VPNs.

The customer carrier can connect to the backbone carrier in the following ways:

- Using IGP and Label Distribution Protocol (LDP)
  - MPLS has to be enabled on the link between the backbone carrier and the customer carrier. LDP is used for label distribution.
  - IGP is used for route exchange.
- Using Multiprotocol BGP (MP-BGP)
  - MP-BGP is used for label and route exchange. There is no need for LDP.

https://t.me/learningnets



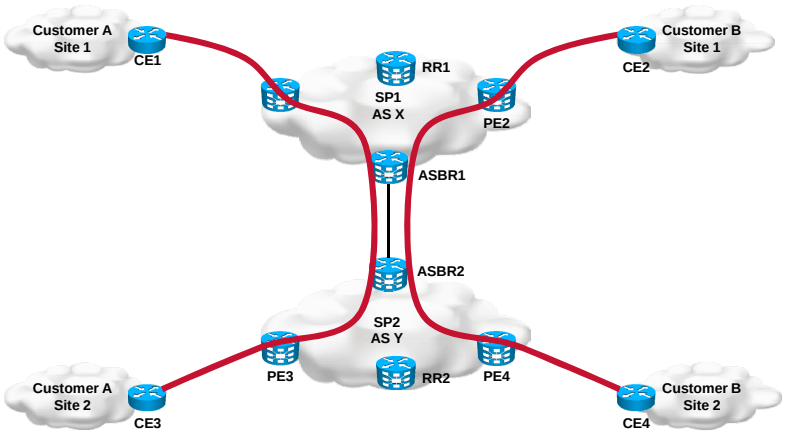
In traditional service provider networks, MPLS VPN was mostly used inside one autonomous system (AS). If customers did not use the same service provider in all branch offices, they would not be able to establish an MPLS VPN.

Inter-AS introduces techniques to establish MPLS VPNs across multiple autonomous systems. Service providers have to establish interconnection, exchange VPN information, and build VPN tunnels.

https://t.me/learningnets

Inter-AS MPLS VPN

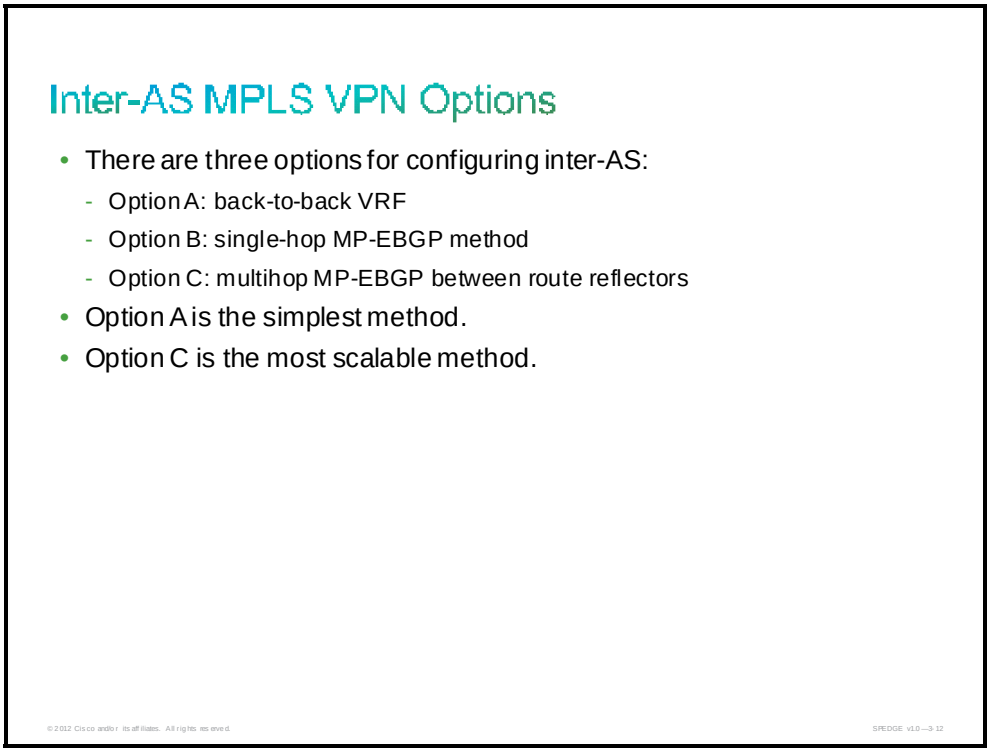
- An MPLS VPN tunnel is established across two service providers.



© 2012 Cisco and/or its affiliates. All rights reserved.

In this example, two customers with two sites are connected to different service providers. An MPLS VPN tunnel is established using an inter-AS connection between service providers.

https://t.me/learningnets



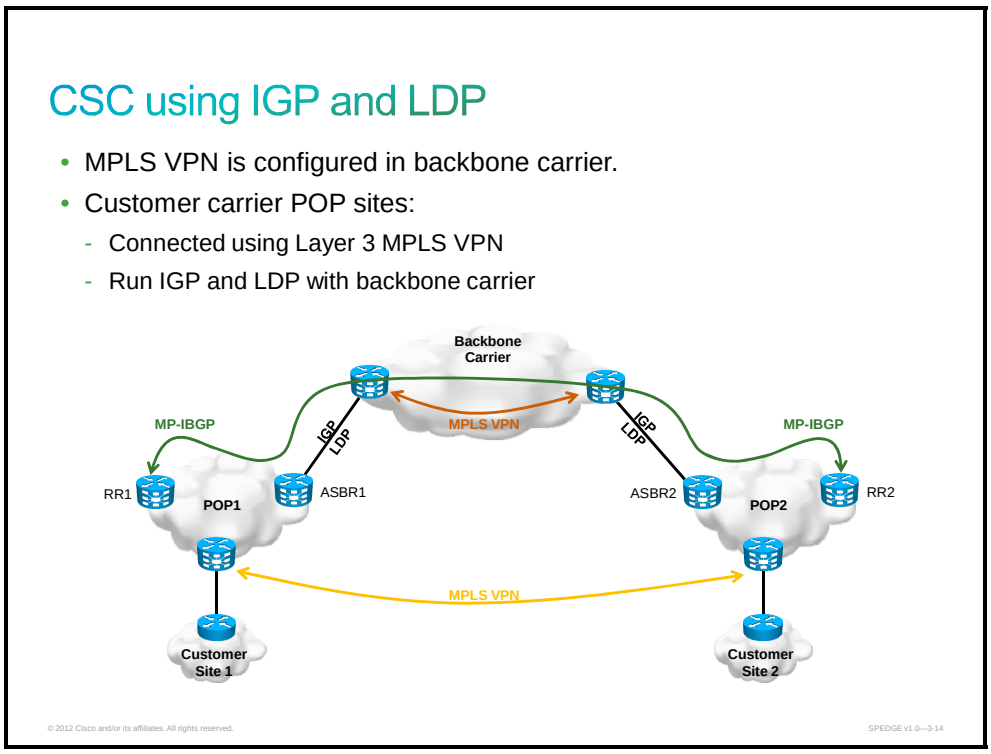
There are three basic techniques for establishing inter-AS:

- Option A: back-to-back VRF
  - The simplest method but not scalable
- Option B: single-hop Multiprotocol Exterior Border Gateway Protocol (MP-EBGP)
  - Scalable method
  - Some routing overhead in Autonomous System Boundary Routers (ASBRs)
- Option C: multihop MP-EBGP between route reflectors
  - Scalable method
  - End-to-end VPN between PE routers

https://t.me/learningnets

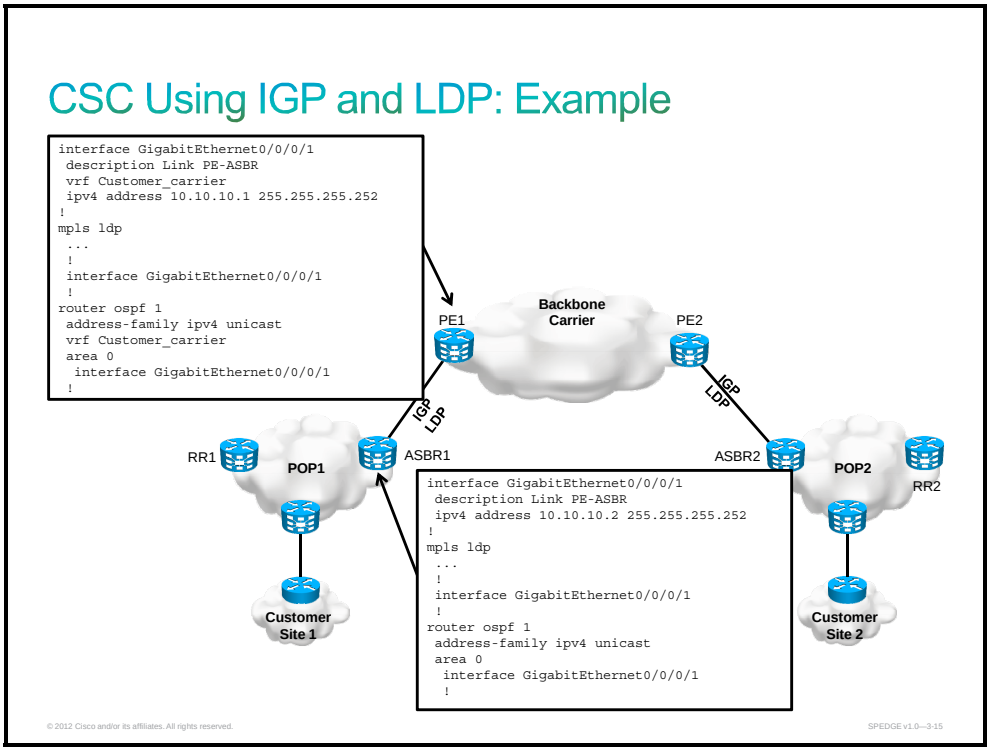
CSC Models

This topic describes different CSC models.



This CSC implementation builds on the MPLS and LDP model. The resulting end-to-end LSP enables the customer carrier to establish a peer relationship between its PE routers that are supporting the end customer. It then enables the customer carrier to use MPLS VPNs to support its end customers over an end-to-end VPN.

https://t.me/learningnets



To configure CSC using IGP and LDP, you have to do the following:

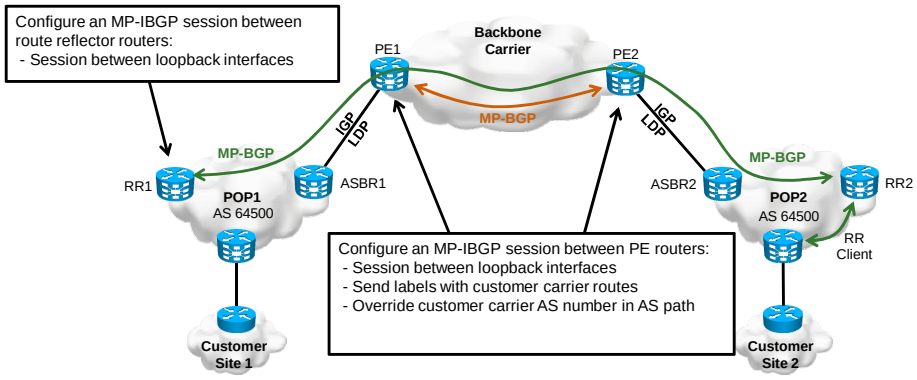
- Enable MPLS LDP for label exchange on a link connecting the backbone carrier and the customer carrier
- Configure IGP (OSPF in this example) to exchange routing information

It is important to have connectivity between route reflectors to establish an MP-BGP session.

https://t.me/learningnets

CSC Using IGP and LDP: Example (Cont.)

- Backbone carrier establishes MPLS VPN for customer carrier.
- Customer carrier establishes MPLS VPN for end customers.



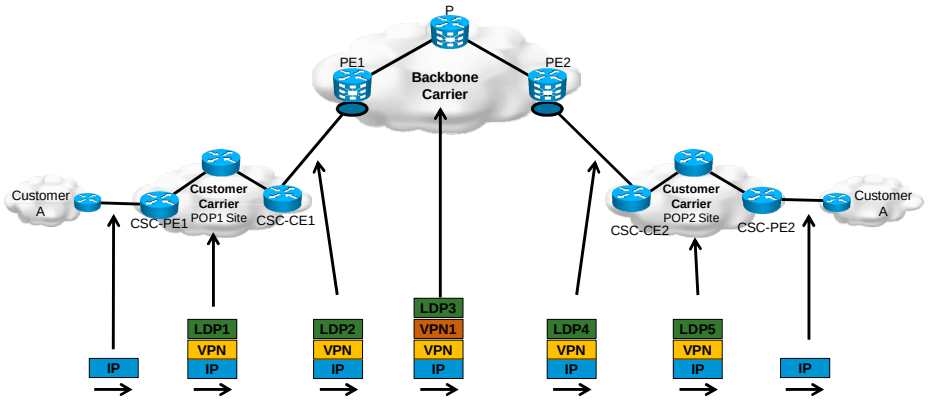
MP-BGP is established between the route reflector routers of the customer carrier. When IBGP is established between two routers, they ignore routes from same AS by default. You can configure the PE router of the backbone carrier to override the AS number so that BGP routes are accepted by IBGP. On routers using Cisco IOS XR Software, you can use the **remove-private-as** command. On routers running Cisco IOS Software, the **as-override** command can be used.

Backbone carriers should be able to send labels with IP prefixes. On IOS XR routers, the **address-family ipv4 labeled-unicast** or **address-family ipv6 labeled-unicast** commands are used. On IOS routers, you have to configure the **send-label** parameter for the BGP peer.

https://t.me/learningnets

Data Flow Example

- When an IP packet enters the customer carrier VPN, an LDP label is attached to it.
- When the packet arrives at the backbone carrier, another VPN label is attached to it.



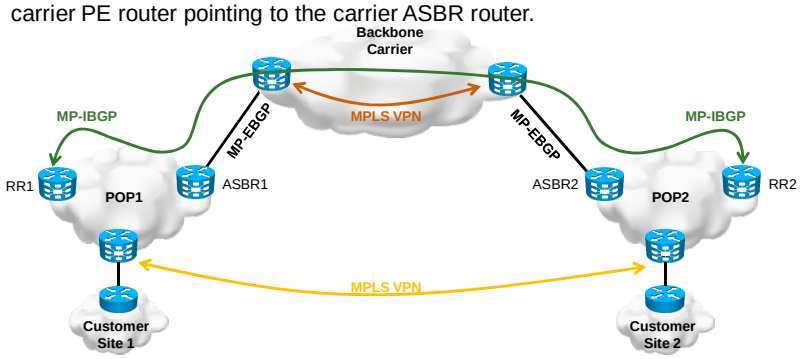
© 2012 Cisco and/or its affiliates. All rights reserved.

In this example, a customer sends an IP packet to the service provider. A VPN label is attached to the packet. This label is used to identify packets from the same VPN. Another LDP label is attached to packet in order to forward the packet in the carrier network. When the packet enters the backbone carrier network, another VPN label is attached to the packet. This label identifies all packets from this provider.

https://t.me/learningnets

CSC Using MP-BGP

- MPLS VPN is configured in backbone carrier.
- Customer carrier POP sites:
  - Connected using Layer 3 MPLS VPN
  - Run MP-EBGP with backbone carrier ASBR
  - Use /32 loopback address for MP-IBGP sessions between route reflectors.
  - On Cisco IOS XR routers, a static route should be configured on the backbone carrier PE router pointing to the carrier ASBR router.



© 2012 Cisco and/or its affiliates. All rights reserved.

When configuring CSC using MP-BGP, LDP is not used for label distribution. MP-BGP can exchange labels and route information between BGP peers.

If you are using routers with Cisco IOS XR Software, you have to configure a new address family using the **ipv4 labeled-unicast address-family** or **ipv6 labeled-unicast address-family** commands.

If you are using routers running Cisco IOS Software, **send-label** should be configured under the BGP peer address family.

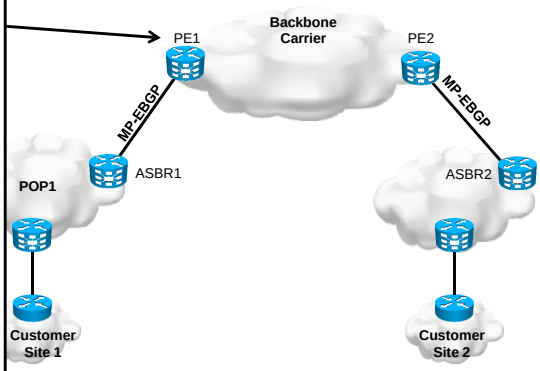
When you configure a BGP session between customer carrier POP sites, you should use /32 loopback addresses for source and destination IP addresses. If the IP address mask is not /32, it can cause problems with label assignment on backbone carrier PE routers.

If you are using a router running IOS XR Software as the backbone carrier PE router, you have to configure a static route toward the carrier ASBR router pointing to the physical interface that connects both routers.

https://t.me/learningnets

CSC Using MP-BGP: Example

```
interface GigabitEthernet0/0/0/1
 description Link PE-ASBR
 vrf Customer_carrier
 ipv4 address 10.10.10.1 255.255.255.252
!
router static
 vrf Customer_carrier
  address-family ipv4 unicast
   10.10.10.2/32 GigabitEthernet0/0/0/1
 !
router bgp 64500
 vrf Customer_carrier
  rd 1:220
  address-family ipv4 unicast
   redistribute connected
   allocate-label all
  !
  neighbor 10.10.10.2
   remote-as 64512
   update-source GigabitEthernet0/0/0/1
   address-family ipv4 unicast
    route-policy pass in
    route-policy pass out
    as-override
    next-hop-self
   !
   address-family ipv4 labeled-unicast
    route-policy pass in
    route-policy pass out
    as-override
    next-hop-self
   !
```

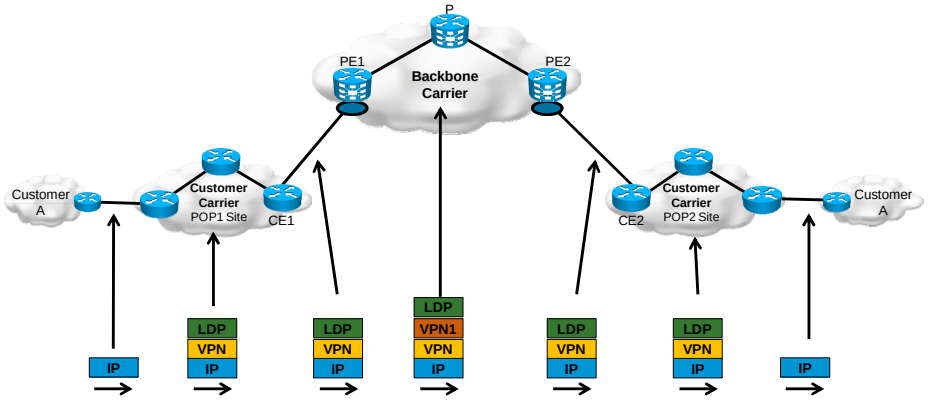


© 2012 Cisco and/or its affiliates. All rights reserved.

This example shows how to configure a BGP peer-facing customer carrier on the PE router of the backbone carrier. The labeled-unicast address family is configured and the **as-override** command is used to rewrite the local AS in the AS path.

Data Flow Example

- When an IP packet enters the customer carrier VPN, an LDP label is attached to it.
- When the packet arrives at the backbone carrier, another VPN label is attached to it.



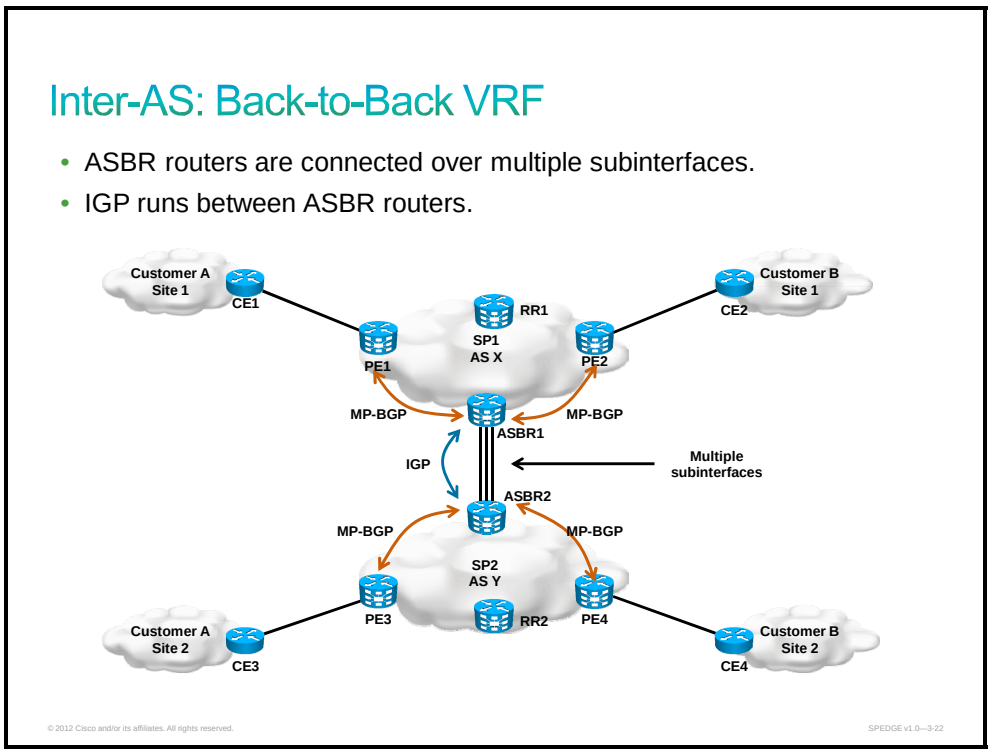
© 2012 Cisco and/or its affiliates. All rights reserved.

Data flow when using MP-BGP is similar to data flow when using LDP and IGP. First, a VPN label is used to identify the customer VPN and a second VPN label is used to identify the customer carrier VPN.

https://t.me/learningnets

Inter-AS

This topic describes the inter-AS method for interconnecting service provider networks.

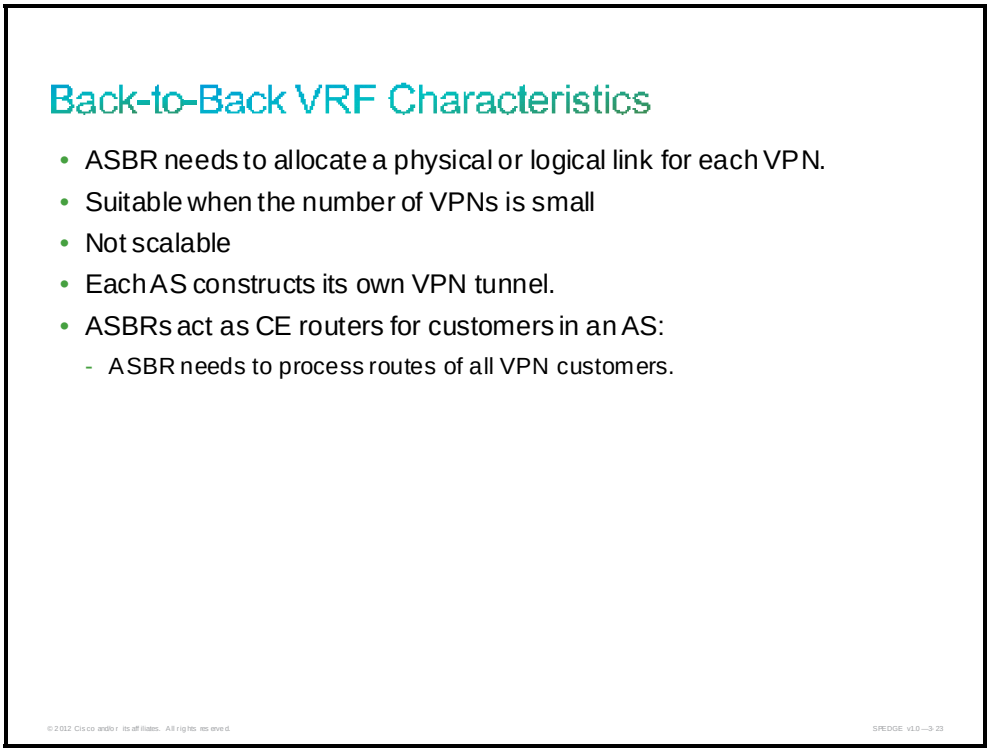


You can configure inter-AS functionality using different techniques. The first is called Option A, or the back-to-back VRF method.

ASBRs are interconnected using multiple interfaces or subinterfaces. Each interface or subinterface is used to carry the traffic of its own VPN.

Each ASBR is acting as a PE router for its customers and a CE router for customers of other service providers. Some IGP is used to exchange customer routing information between ASBRs.

https://t.me/learningnets

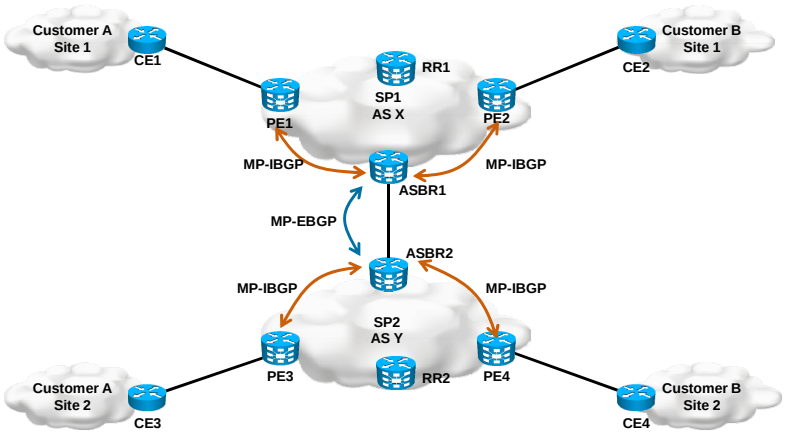


The back-to-back VRF inter-AS method is only suitable when two service providers have a small number of VPN tunnels. This method is easy to configure, but it is not scalable. Because ASBRs act as CE routers for all VPN customers, these routers have to maintain routes from all customers in their memory.

https://t.me/learningnets

Inter-AS: Single-Hop MP-EBGP Method

- BGP is used to signal VPN labels between the AS boundary routers.
- Higher scalability

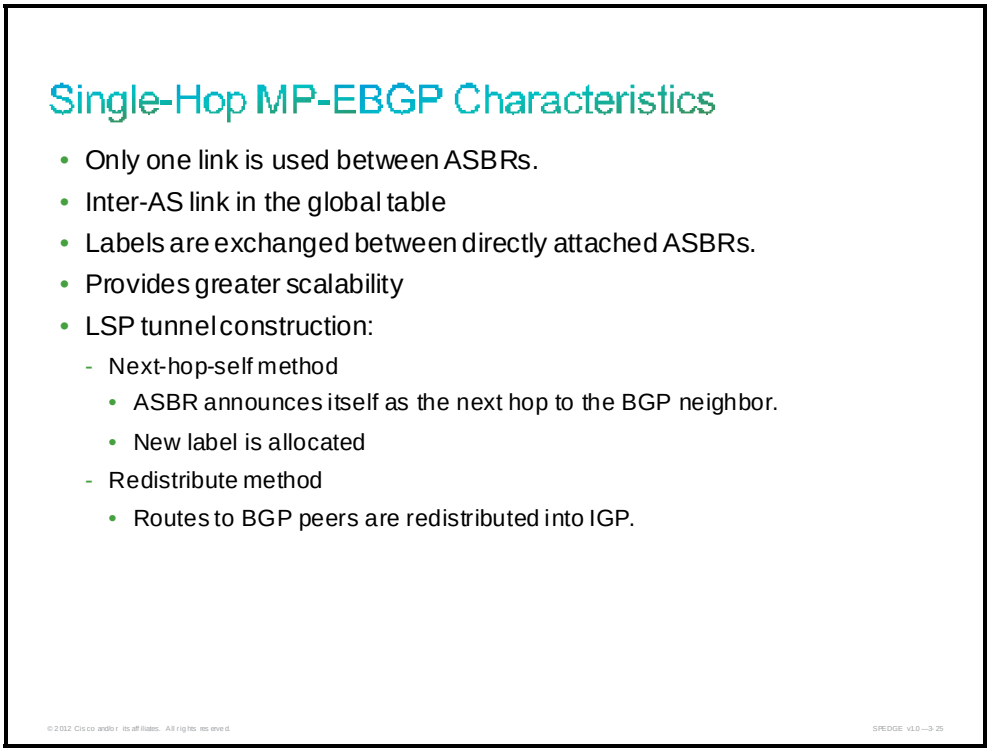


© 2012 Cisco and/or its affiliates. All rights reserved.

The second method to configure inter-AS is called Option B, or the single-hop MP-EBGP method.

Inside an AS, normal MPLS and BGP are used to transfer VPN information and construct the LSP tunnel. Between autonomous systems, the single-hop MP-EBGP method is used to transfer VPN information and construct the LSP tunnel.

https://t.me/learningnets



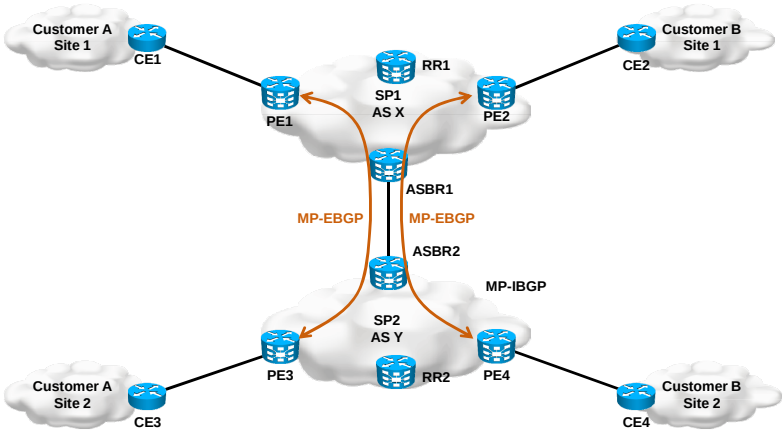
The single-hop MP-EBGP method provides more scalability than the back-to-back VRF method. Only one link is configured between service providers. MP-BGP is used co exchange routing and label information between directly connected routers.

To construct the LSP path, next-hop addresses should be reachable. Routes to BGP peers can be redistributed to the provider IGP, or next-hop-self can be used by having the ASBR replace the next hop with its IP address.

https://t.me/learningnets

Inter-AS: Multihop MP-EBGP

- Labeled IPv4 routes are redistributed by EBGP between neighboring autonomous systems.
- BGP is used for label distribution.

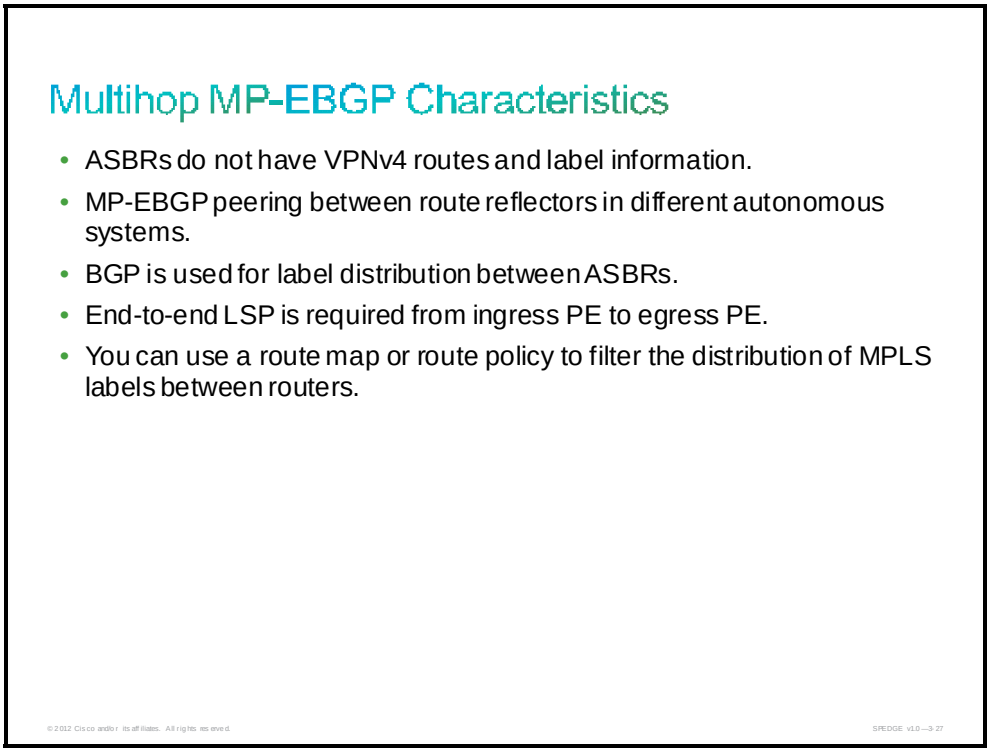


© 2012 Cisco and/or its affiliates. All rights reserved.

The third method to configure inter-AS is called Option C, or multihop MP-EBGP.

Because BGP only requires the TCP connection to form a BGP neighbor and transfer route information, this third method transfers VPN route information between the source and destination PEs directly over multihop MP-EBGP, and constructs a public network LSP tunnel between the source and destination PEs. When there is route reflector in service, a provider network BGP session can be established between route reflectors.

https://t.me/learningnets



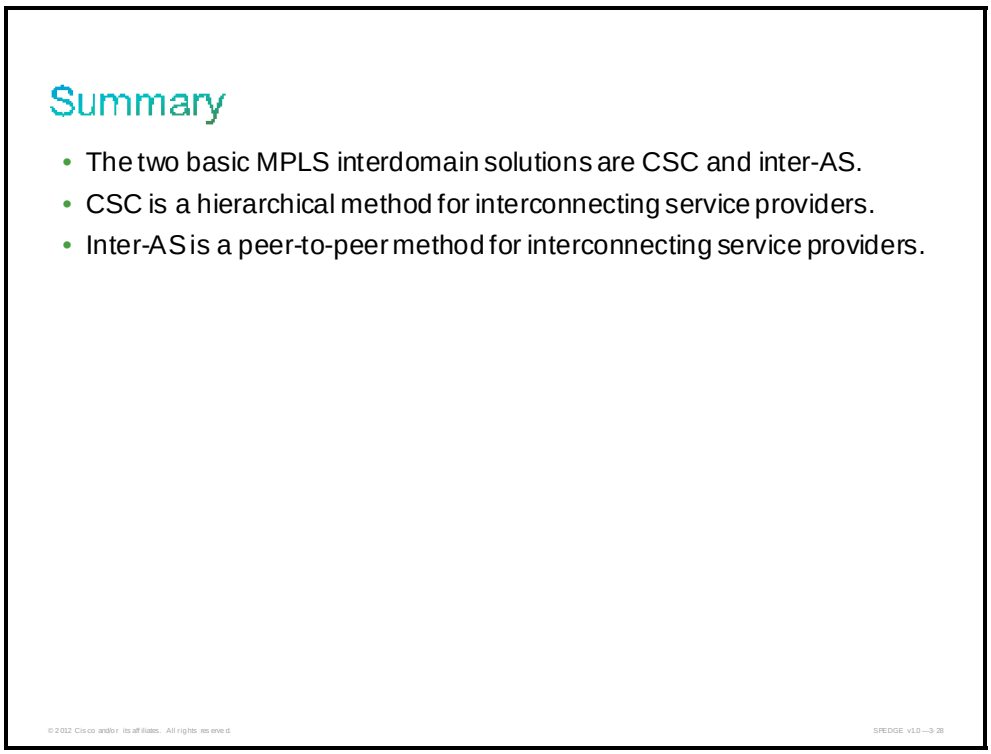
To use the multihop MP-EBGP method, end-to-end LSP is required from the ingress PE (or route reflector) to the egress PE (or route reflector). This method is highly scalable, because there is no route overhead on ASBRs.

You can use a route map or route policy to filter the distribution of MPLS labels between routers.

https://t.me/learningnets

Summary

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



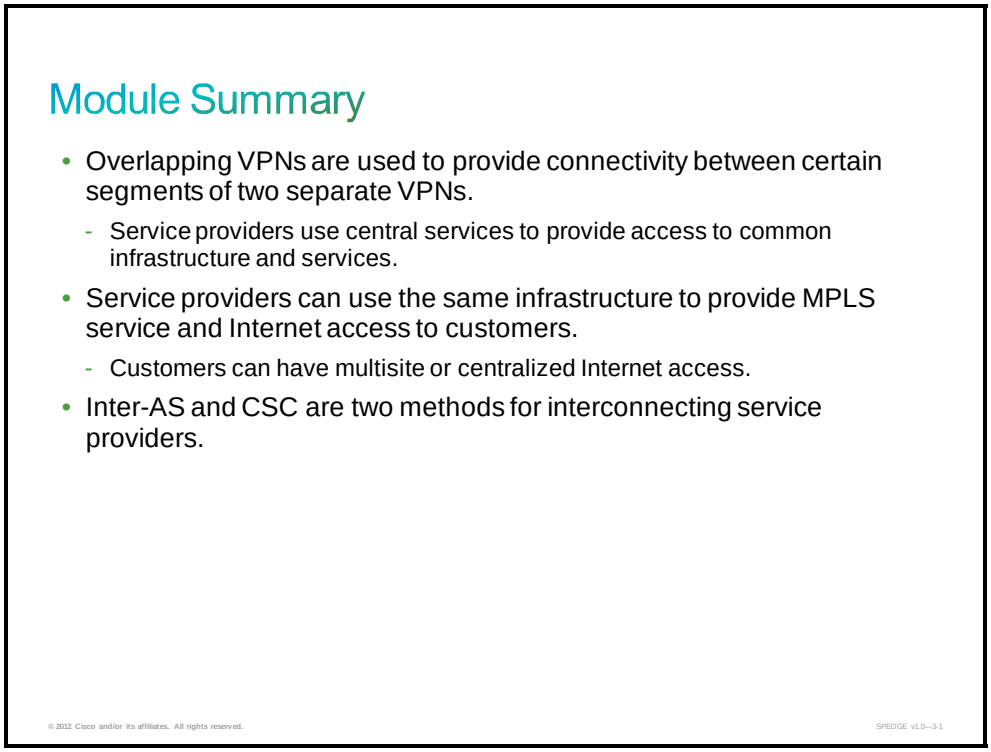
https://t.me/learningnets

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Module Summary

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



https://t.me/learningnets

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Module Self-Check

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

Q1) Why do you need a selective VRF import command? (Source: Implementing Complex MPLS Layer 3 VPNs)

Q2) Why do you need a selective VRF export command? (Source: Implementing Complex MPLS Layer 3 VPNs)

Q3) Who are the typical users of overlapping VPNs? (Source: Implementing Complex MPLS Layer 3 VPNs)

Q4) What are the connectivity requirements for overlapping VPNs? (Source: Implementing Complex MPLS Layer 3 VPNs)

Q5) What are the typical usages for a central services VPN topology? (Source: Implementing Complex MPLS Layer 3 VPNs

Q6) Why do you need the managed CE routers service? (Source: Implementing Complex MPLS Layer 3 VPNs)

Q7) What is the main difference between the managed CE routers service and the typical central services VPN topology? (Source: Implementing Complex MPLS Layer 3 VPNs)

https://t.me/learningnets

Q8) What are the two major design models for implementing Internet access and MPLS VPNs? (Choose two.) (Source: Implementing Internet Access and MPLS Layer 3 VPNs)

A) using a separate VPN for Internet access
B) using global route leaking
C) using global routing on PE routers
D) using default routes

Q9) What is the recommended implementation option for using global routing to provide Internet access? (Source: Implementing Internet Access and MPLS Layer 3 VPNs)

A) The separate subinterfaces are placed in separate VRFs.
B) The Internet subinterface is not placed in a VRF.
C) One interface for the Internet and the MPLS VPN is recommended.
D) Static routes are always needed for Internet access.

Q10) Which two are drawbacks of Internet access that is implemented as a separate VPN? (Choose two.) (Source: Implementing Internet Access and MPLS Layer 3 VPNs)

A) Full Internet routing cannot be carried inside a VRF.
B) The customer cannot receive full Internet routes.
C) Default routing toward the Internet gateways must be used, potentially resulting in suboptimal routing.
D) The customer must use default routes.

Q11) Which two of these are true about the CSC model? (Choose two.) (Source: Introducing MPLS Interdomain Solutions)

A) CSC is a peer-to-peer model.
B) CSC is a hierarchical model.
C) Multiple customer carriers can be connected to a single CSC backbone.
D) End customer routes are inserted in the CSC backbone.

Q12) Which two protocols are used to exchange label and route information between the backbone and customer carriers? (Choose two.) (Source: Introducing MPLS Interdomain Solutions)

A) CDP
B) OSPF and LDP
C) MP-BGP
D) NTP

Q13) Which two of these are true about the inter-AS model? (Choose two.) (Source: Introducing MPLS Interdomain Solutions)

A) Inter-AS is a peer-to-peer model.
B) Inter-AS is a hierarchical model.
C) The back-to-back VRF method is suitable for small environments but is not a scalable solution.
D) In the multihop EBGP method, the ASBR router contains customer routes.

Q14) What are the names of the three basic options for configuring inter-AS? (Source: Introducing MPLS Interdomain Solutions)

_____

_____

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Module Self-Check Answer Key

Q1) A selective VRF import command allows you to select routes to be imported into a VRF based on criteria other than the VRF RT.

Q2) A selective VRF export command allows you to attach specific RTs to a subset of routes exported from a VRF. (By default, the same RTs get attached to all exported routes.)

Q3) Companies that use MPLS VPNs to implement both intranet and extranet services, or a security-conscious company that wants to limit visibility between different departments in the organization

Q4) Selected sites in a VPN can communicate only with sites within their VPN. Other selected sites can communicate with sites in their VPN and selected sites in a second VPN.

Q5) In solutions where some sites (server sites) can communicate with all other sites, but all the other sites (client sites) can communicate only with the server sites

Q6) If the service provider is managing the customer routers, it is convenient to have a central point that has access to all CE routers but not to the other destinations at customer sites.

Q7) The VRF and RD design is similar to that of a central services VPN. The managed CE routers service combines a service VPN and simple VPN topology like the central services VPN. However, the route export statement uses a route policy to limit the exported addresses to the loopback address of the managed routers.

Q8) A, C

Q9) B

Q10) A, C

Q11) B, C

Q12) B, C

Q13) A, C

Q14) Option A, or back-to-back VRF; Option B, or the single-hop MP-EBGP method; and Option C, or the multihop MP-EBGP method

https://t.me/learningnets

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Module 4

Layer 2 VPNs and Ethernet Services

Overview

Virtual Private Wire Service (VPWS) enables geographically separated sites to be interconnected over a virtual point-to-point Layer 2 circuit. The virtual connection can cross a Multiprotocol Label Switching (MPLS) or IP network. Virtual Private LAN Services (VPLS) enable remote LAN segments to be linked as a single bridged domain over an MPLS network. The full functions of the traditional LAN, such as MAC address learning, aging, and switching, are emulated across all the remotely connected LAN segments that are part of a single bridged domain.

A service provider can offer VPLS service to multiple customers over the MPLS network by defining different bridged domains for different customers. Packets from one bridged domain are never carried over or delivered to another bridged domain, thus ensuring the privacy of the LAN service. VPLS transports Ethernet 802.3, VLAN 802.1Q, and VLAN-in-VLAN (QinQ) traffic across multiple sites that belong to the same Layer 2 broadcast domain. VPLS offers simple VLAN services that include flooding broadcast, multicast, and unknown unicast frames that are received on a bridge. The VPLS solution requires a full mesh of pseudowires (PWs) that are established among provider edge (PE) routers. The VPLS implementation is based on Label Distribution Protocol (LDP)-based PW signaling.

This module describes the VPWS and VPLS technology and implementation on Cisco platforms.

Module Objectives

Upon completing this module, you will be able to describe Layer 2 VPNs and Ethernet services. You will be able to meet these objectives:

- Describe Layer 2 VPNs that are available in the MPLS and IP core
- Describe AToM
- Describe Ethernet services that are used in the service provider network

https://t.me/learningnets

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 1

Introducing Layer 2 VPNs

Overview

There is an ever-increasing demand for the transport of Layer 2 and Layer 3 over a common backbone. This lesson introduces the Virtual Private Wire Service (VPWS) and Virtual Private LAN Services (VPLS). VPWS offers point-to-point virtual connections, while VPLS provides LAN-similar multipoint connectivity. A subset of VPWS is Any Transport over MPLS (AToM). AToM allows a Multiprotocol Label Switching (MPLS) network to provide end-to-end transport for Layer 2 frames and cells. It provides support for Ethernet, PPP, High-Level Data Link Control (HDLC), Frame Relay, and ATM.

Objectives

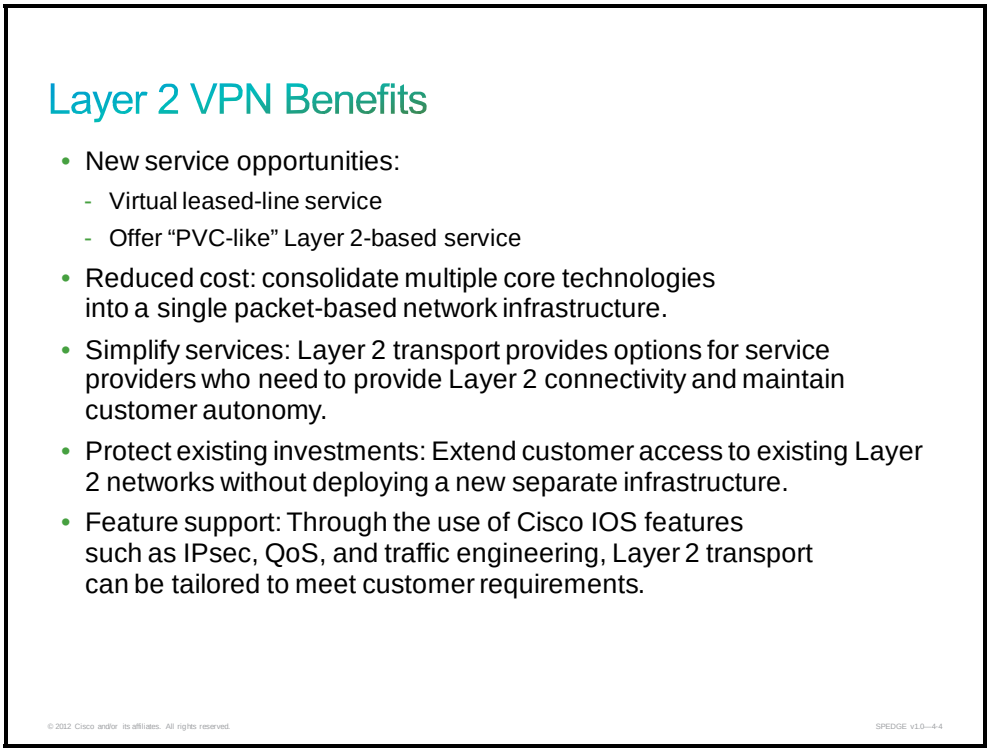
Upon completing this lesson, you will be able to describe Layer 2 VPNs that are available in the MPLS and IP core. You will be able to meet this objective:

- Explain Layer 2 VPN services that are available with IP and MPLS core

https://t.me/learningnets

Layer 2 VPN Overview

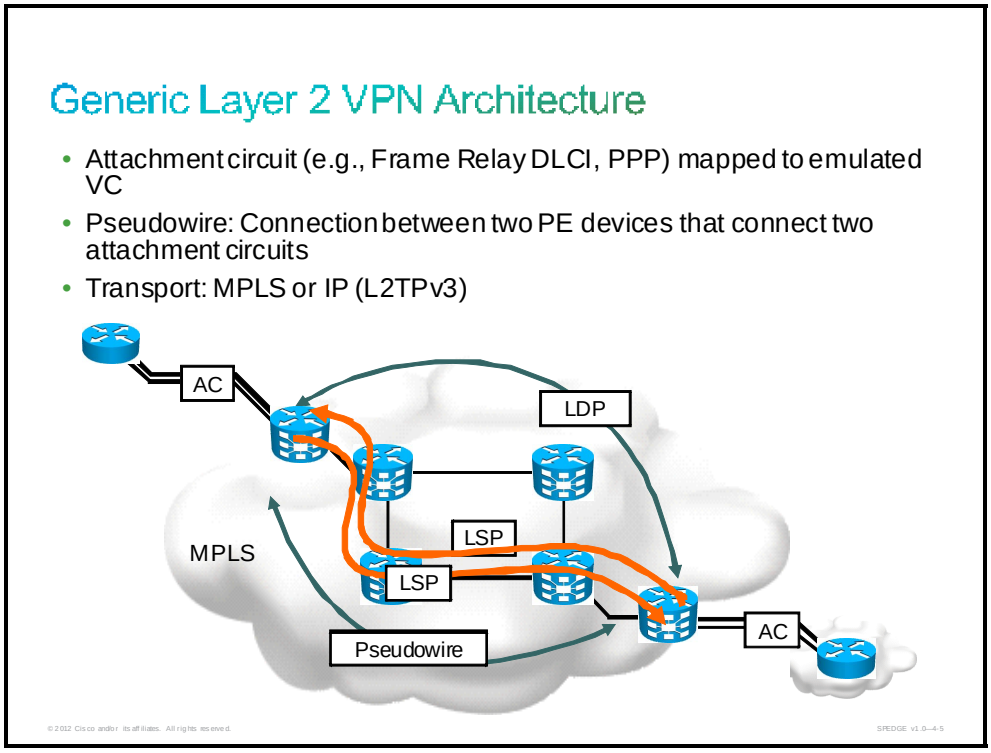
This topic explains Layer 2 VPN services that are available with IP and MPLS core.



Layer 2 VPN technologies offer a range of benefits to service providers and enterprises:

- **New service opportunities:** The virtual leased lines offer connectivity service that resembles the traditional costly permanent virtual circuits in Frame Relay or ATM environments.
- **Cost savings:** The consolidation of multiple core technologies into a single packet-based network infrastructure lowers the overall cost of system installation and maintenance.
- **Simple connectivity model:** Layer 2 transport provides options for service providers that need to provide Layer 2 connectivity and maintain customer autonomy.
- **Investment protection:** Service providers can extend customer access to existing Layer 2 networks without deploying a new separate infrastructure.
- **Feature support:** Through the use of Cisco IOS and IOS XR features such as IPsec, quality of service (QoS), and traffic engineering, Layer 2 transport can be tailored to meet customer requirements.

https://t.me/learningnets



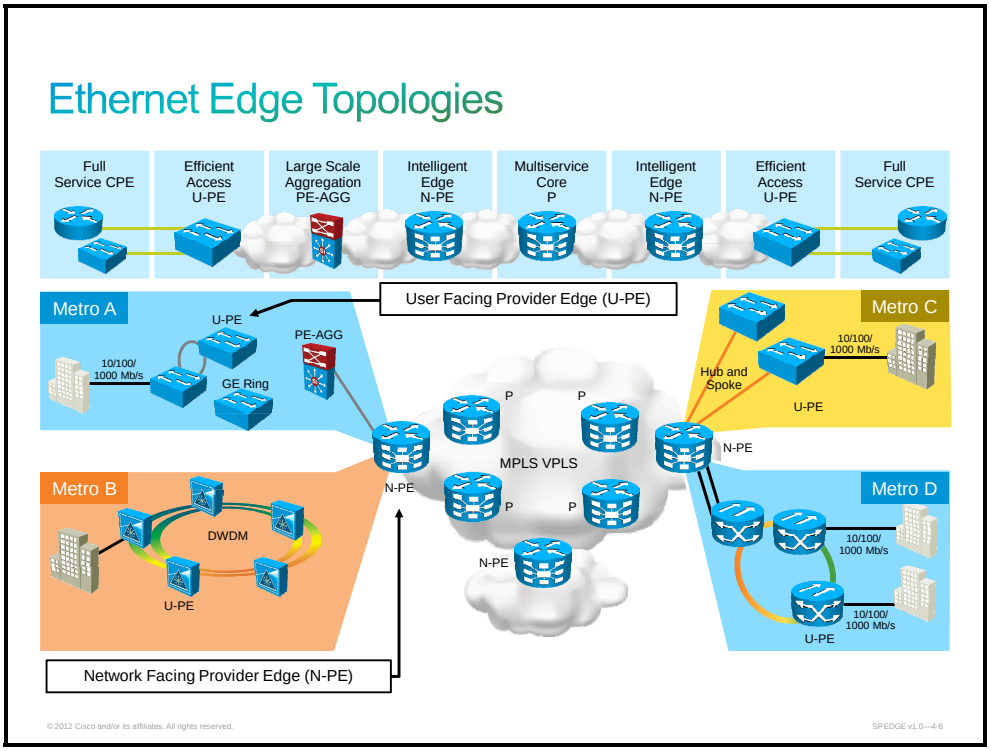
The Layer 2 VPN solution encompasses three main elements:

- An attachment circuit is the circuit or link directly connected to the provider edge (PE) system that is virtually extended to the other end of the provider cloud. It can represent a Frame Relay data link connection identifier (DLCI), PPP and HDLC link, ATM virtual path identifier (VPI) and virtual channel identifier (VCI) pair, Ethernet port, or VLAN. The attachment circuit is mapped to the emulated virtual circuit (VC) for transport through the service provider core.
- Pseudowire (PW) is a point-to-point connection over a packet-switching network. The PW emulates the operation of a “transparent wire” by linking two distant attachment circuits attached to two different PE routers.
- Transport infrastructure is the third element. The transport network can either be MPLS-enabled, and thus capable of supporting any Layer 2 topology (point-to-point, point-to-multipoint, or multipoint-to-multipoint), or IP-based. In the IP core, the Layer 2 Tunneling Protocol (L2TP) protocol supports point-to-point connections only.

Frames are received on an ingress interface by the ingress PE router. At this point, the frame is a raw Layer 2 frame. In the case of MPLS transport, the ingress PE router encapsulates it into MPLS and tunnels it across the backbone to the egress PE router. The egress PE router decapsulates the packet and reproduces the raw Layer 2 frame on the egress interface.

The frames are carried across the MPLS backbone using a label stack of two labels. The top label is used to propagate the packet from the ingress PE router to the correct egress PE router. The second label is used by the egress PE router to forward out the packet on the correct interface. This process is somewhat similar to an MPLS VPN, where the egress PE router uses a VPN label. The top label is called the tunnel label. This name indicates that its use is to tunnel the packet across the MPLS backbone to the egress PE router. The second label is called the VC label. The name indicates that its use is to map the packet to an outgoing VC or link.

https://t.me/learningnets

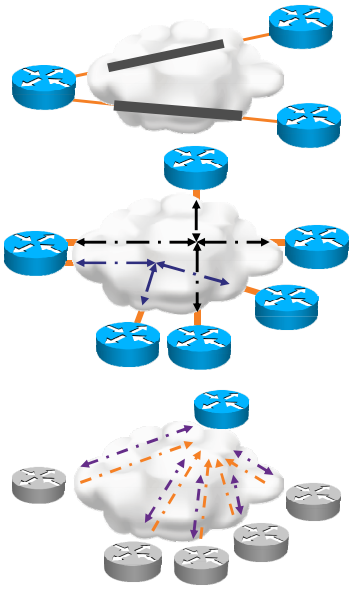


Carrier Ethernet environments are typically designed and deployed according to the architecture seen in the figure. It illustrates the various layers: the multiservice core, the network edge with intelligent services, aggregation of large amounts of access circuits, the access layer, and the customer equipment. In many cases, a specific carrier or Metro Ethernet solution may not contain all of these layers. In fact, in some cases the architectural functions can be merged into a single layer. For example, various combinations of network technologies and topologies can be formed to deliver Ethernet services without passing through a core network. In this context, these network technology and topology combinations can be viewed as separate from the interconnecting core network, and are hence referred to as Metro Ethernet islands (or simply islands).

https://t.me/learningnets

Metro Ethernet Forum Service Types

- E-Line:
  - Ethernet private lines
  - Ethernet virtual private lines
  - Ethernet Internet access
- E-LAN:
  - Multipoint Layer 2 VPNs
  - Transparent LAN service
  - Foundation for IPTV and multicast networks
- E-Tree:
  - Also known as rooted multipoint
  - Leaves can communicate with one or more roots
  - Leaves do not communicate with other leaves
  - Targeted at multihost separation
  - Enabler for mobile backhaul and triple-play infrastructure



© 2012 Cisco and/or its affiliates. All rights reserved.

To bring about the industry acceptance of Ethernet services, it was necessary to clarify and standardize the service range. Recognizing this requirement, the industry created the Metro Ethernet Forum (MEF), which played a key role in defining the three major services:

- **E-Line**: a service connecting two customer Ethernet ports over a WAN. This service is further subdivided into Ethernet private lines (EPLs), Ethernet virtual private lines (EVPLs), and Ethernet Internet access (EIA).
- **E-LAN**: a multipoint service connecting a set of customer endpoints, giving the appearance to the customer of a bridged Ethernet network connecting the sites. This transparent LAN service is often referred to as a multipoint Layer 2 VPN. It lays the foundation for IPTV and multicasting applications.
- **E-Tree**: a multipoint service connecting one or more roots and a set of leaves, but preventing interleaf communication. This service is also known as rooted multipoint. Specifically, the leaves can communicate with one or more roots, but not with other leaves. The service provides an ideal mechanism for multihost separation. It is considered a major enabler for mobile backhaul and triple-play infrastructure.

All these services provide standard definitions of such characteristics as bandwidth, resilience, and service multiplexing, allowing customers to compare service offerings and facilitating service level agreements (SLAs).

https://t.me/learningnets

MEF Services Standards Map

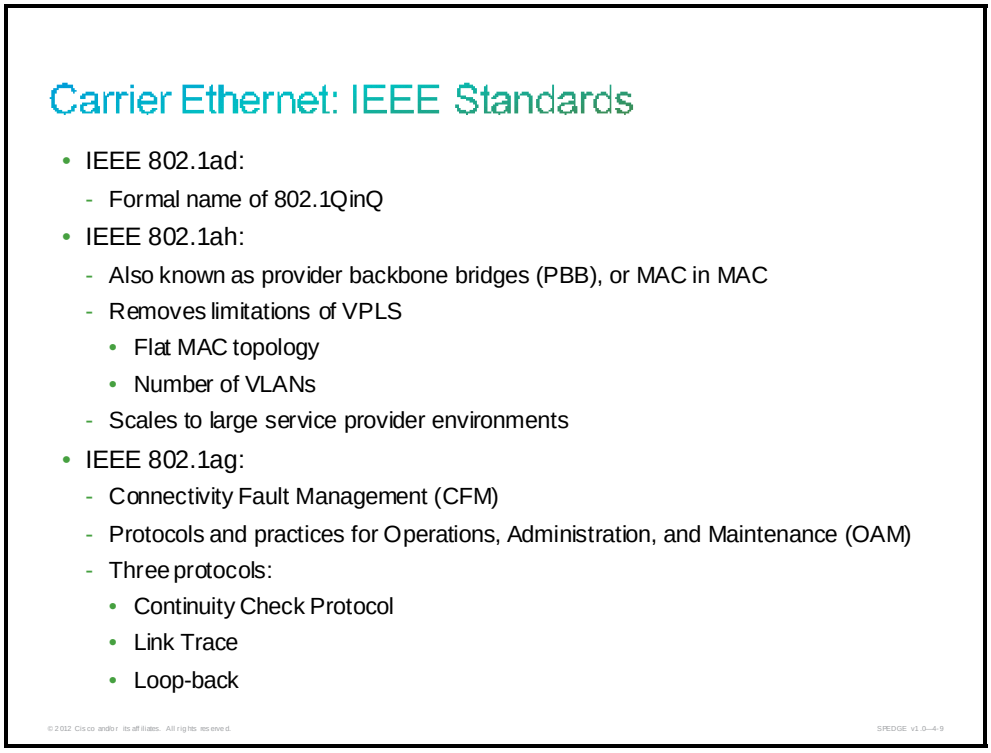
Metro Ethernet Forum		IETF (MPLS)	IEEE	Cisco service name
E-Line (point-to-point)	Ethernet private line (EPL)	Virtual Private Wire Service (VPWS)	QinQ, .1ad	Ethernet Wire Service (EWS)
	Ethernet virtual private line (EVPL)		.1Q	Ethernet Relay Service (ERS)
E-LAN (multipoint)	Transparent LAN Service (TLS)	Virtual Private LAN Service (VPLS)	QinQ, .1ad	Ethernet Multipoint Service (EMS)
	Ethernet Virtual Connection Service (EVCS)		.1Q	Ethernet Relay Multipoint Service (ERMS)

© 2012 Cisco and/or its affiliates. All rights reserved.

MEF does not define only the three major categories (E-Line, E-LAN, and E-Tree), but many variants of them. One classification of the variants is based on the Ethernet virtual circuit (EVC) mode: port- or VLAN-based.

In addition, the terms E-line, E-LAN, and E-Tree are not the only ones used in the industry. IETF refers to these same services as VPWS and VPLS. The IETF naming is used predominantly throughout this course.

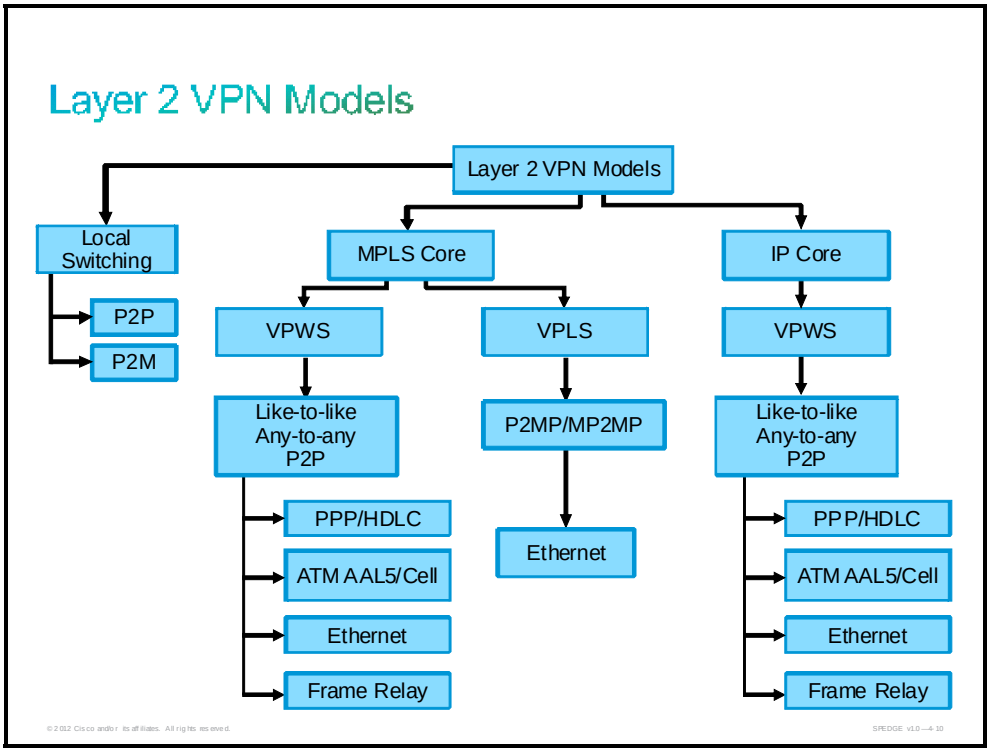
https://t.me/learningnets



Many standardization bodies have issued standards that define basic functions and enhancements of the Carrier Ethernet architecture. IEEE plays a major role in this field and contributed many important standards, such as the following:

- IEEE 802.1ad, which is the formal name of the 802.1QinQ standard that allows Ethernet frame encapsulation in multiple VLAN tags.
- IEEE 802.1ah, also known as provider backbone bridges (PBBs), or MAC in MAC. This technology addresses the limitations of VPLS, such as flat MAC topology or limited number of VLANs. It scales to large service provider environments.
- IEEE 802.1ag, also referred to as Connectivity Fault Management (CFM). This set of recommendations defines protocols and practices for Ethernet Operations, Administration, and Maintenance (OAM). Specifically, it comprises three protocols: Continuity Check Protocol, Link Trace, and Loop-back.

https://t.me/learningnets



Layer 2 VPNs are grouped into three main categories: local switching that serves directly connected links, methods for transport in the MPLS core, and transport solutions for IP core.

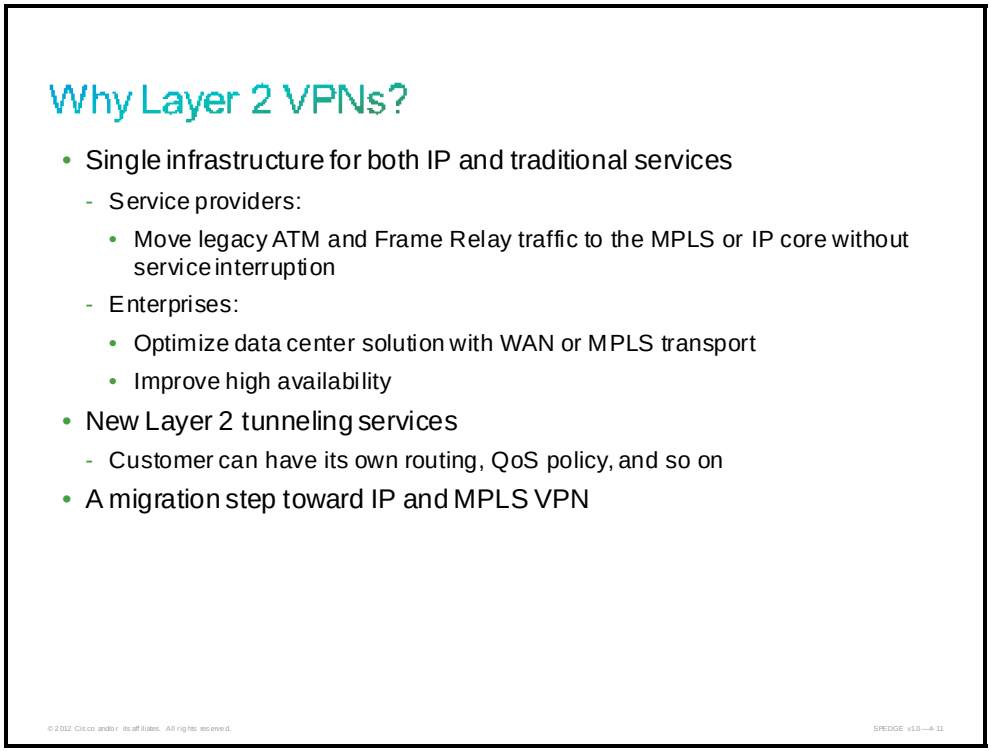
MPLS core transport is further subdivided into VPWS and VPLS).

VPWS is a point-to-point technology. MPLS-based VPWS is called Any Transport over MPLS (AToM) and supports connections between the same interface types (like-to-like), and between different interface types (any-to-any). The supported attached interface types include Ethernet, Frame Relay, and ATM, including ATM adaptation layer 5 (AAL5), PPP, and HDLC.

VPLS offers point-to-multipoint and multipoint-to-multipoint connectivity. It leverages MPLS as the transport infrastructure.

IP core transport uses the Layer 2 Tunneling Protocol version 3 (L2TPv3), and supports only point-to-point connections. The available encapsulations include Ethernet, Frame Relay, and ATM, including AAL5, PPP, and HDLC. They can be linked in any-to-any fashion.

https://t.me/learningnets



Layer 2 VPN allows both service providers and enterprise to build a single infrastructure for both IP and traditional services. The service providers can migrate their existing ATM and Frame Relay traffic to MPLS or IP core without interrupting the customer service. The enterprises can leverage the extended Layer 2 domain to optimize the data center solution. The MPLS transport enables a host of additional high availability extensions.

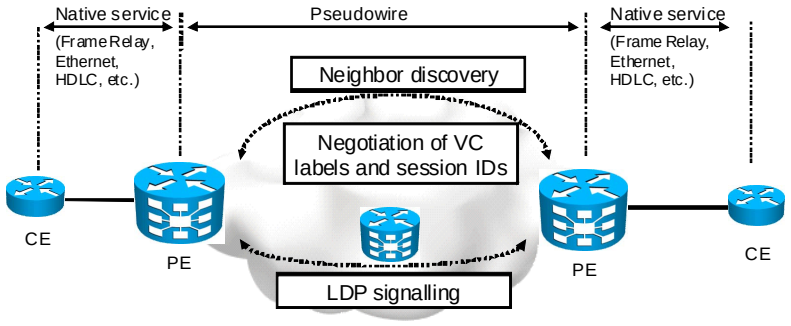
The Layer 2 VPN is a tunneling technology that provides logical separation between the customer and service provider domains, including segmentation of routing, QoS policy, and others.

Layer 2 VPNs also represent an easy-to-implement migration step toward a Layer 3 IP and MPLS VPN system.

https://t.me/learningnets

Layer 2 VPN Control Plane

- Appearance of CE-to-CE native service (transparent service provider network)
- Negotiation of VC labels or session IDs
- Signaling and interworking with native services (e.g., Frame Relay LMI)
- Discovery of other PE VPN members (MP-BGP, LDP)



© 2012 Cisco and/or its affiliates. All rights reserved.

A protocol is required between the PE routers so that they can exchange the VC information.

In the case of MPLS transport, the Label Distribution Protocol (LDP) is used for this purpose. A directed multihop LDP session is established between the PE routers. The egress PE router sends an LDP message in which it indicates the label value to use for a virtual circuit forwarding equivalence class (VC FEC). That label value is then used by the ingress PE router as the second label in the label stack that is imposed to the frames of the indicated VC FEC.

In the case of IP-based transport, the L2TPv3 session exchanges session parameters and not labels.

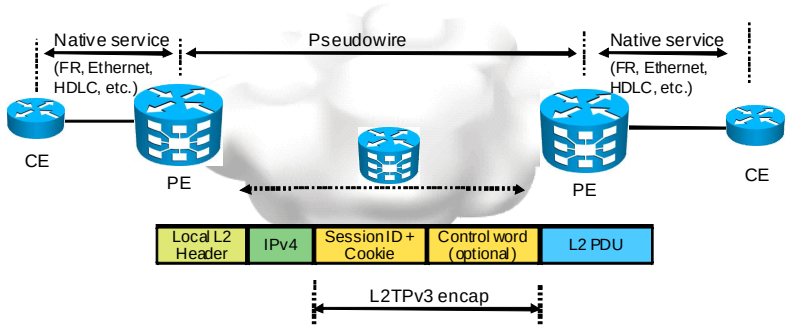
The figure shows a directed multihop LDP session between the ingress and egress PE routers that is used to exchange the VC label. Any ingress-egress PE router pair will need such an LDP session.

The control session is also responsible for providing interworking capabilities with native services, such as Frame Relay Local Management Interface (LMI). The neighbors can be either statically configured or auto-discovered using Multiprotocol Border Gateway Protocol (MP-BGP) extensions or LDP.

https://t.me/learningnets

Layer 2 VPN Data Plane: IP Core

- Transport header: IPv4
- Tunnel header:
  - 4-byte session ID with optional 8-byte cookie
  - Signaled or statically configured
- Layer 2 PDU: Layer 2 specific sublayer with payload (CE Layer 2 PDU)



© 2012 Cisco and/or its affiliates. All rights reserved.

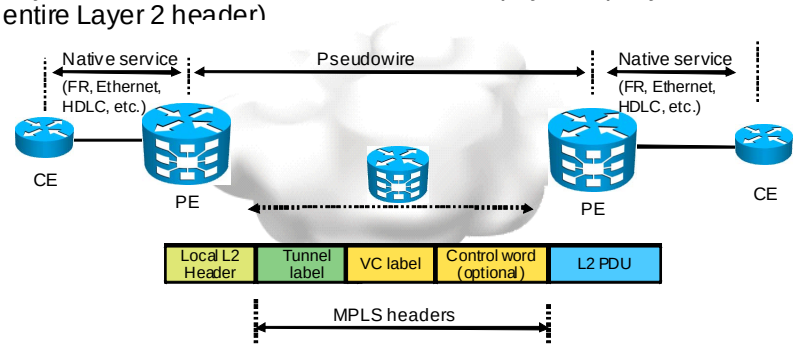
L2TPv3 is used to transport Layer 2 frames over pure IP networks. The entire L2TP packet, including payload and L2TP header, is sent within a UDP datagram. Traditionally, L2TP has been used to carry PPP sessions within an L2TP tunnel. L2TPv3 provides additional security features, improved encapsulation, and the ability to carry data links other than just PPP over an IP network (for example, Frame Relay, Ethernet, ATM, and others).

L2TP overhead includes the transport IP header (20 bytes) and an L2TP header of variable length. The only mandatory field in the L2TP header is the session ID (4 bytes). Optional fields are cookie (8 bytes) and control word. The payload of the L2TPv3 packet is the original Layer 2 protocol data unit (PDU).

https://t.me/learningnets

Layer 2 VPN Data Plane: MPLS Core

- Transport header: MPLS label (4-byte)
  - Signaled via LDP or MPLS-TE
  - Unidirectional path to egress PE
- Tunnel header: VC Label
  - Signaled via directed LDP
- Layer 2 PDU: Control word with customer payload (may not include entire Layer 2 header)

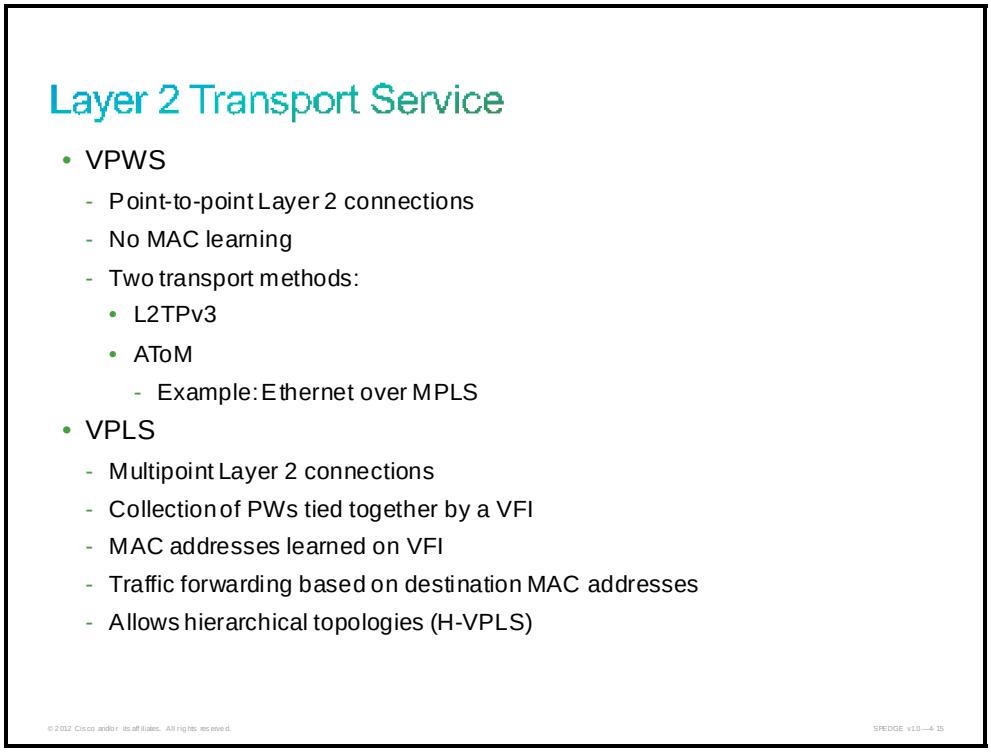


© 2012 Cisco and/or its affiliates. All rights reserved.

MPLS transport is based on the same mechanism that you examined for Layer 3 MPLS VPNs. Hop-by-hop LDP signals a unidirectional path to the egress PE router. A directed LDP session exchanges the VC label that serves as the inner label in the MPLS packet. In general, the use of control words is optional, and configurable. Some transport types, such as Frame Relay Layer 2 VPN, require the use of control words.

Depending on the frame encapsulation, some fields of the original frame, such as checksums, may be stripped before MPLS encapsulation.

https://t.me/learningnets



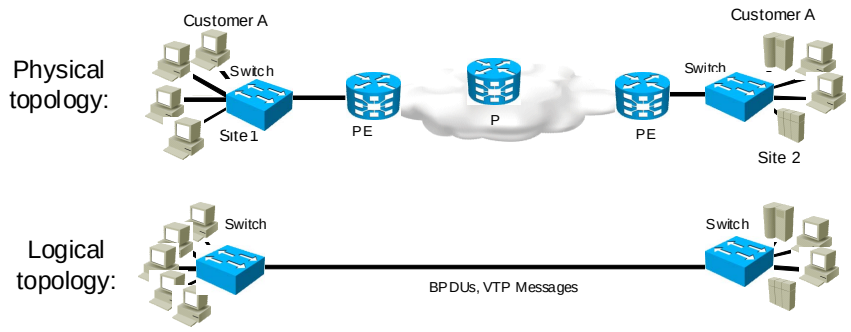
Layer 2 transport services are categorized into two major classes:

- **VPWS:** This service type is used to deploy point-to-point Layer 2 connections. It does not involve MAC learning capabilities. It encompasses two transport methods:
  - L2TPv3
  - AToM. The most common example of AToM is Ethernet over MPLS (EoMPLS).
- **VPLS.** This technology supports multipoint Layer 2 connections by grouping a collection of PWs terminated on a PE router in a Virtual Forwarding Interface (VFI). The VFI represents a virtual extension of the physical circuit attached to the PE system. The VFI resembles a switch that is capable of learning MAC addresses and forwards traffic based on its MAC address table. A VPLS can connect thousands of PEs into a single VLAN and is therefore subject to scalability constraints. To improve the scalability of the solution, hierarchical VPLS (H-VPLS) topologies enable a two-tier deployment of the PE devices.

https://t.me/learningnets

AToM Example: EoMPLS

- Flavor of AToM (attachment circuit: Ethernet; transport: MPLS)
- Attachment circuit can be based on:
  - Port (VC label type 0x0005)
  - VLAN (VC label type 0x0004)
  - Ethernet flow point



© 2012 Cisco and/or its affiliates. All rights reserved.

EoMPLS is the most ubiquitous example of AToM, which transports Ethernet frames over an MPLS infrastructure. The virtual link bridging the Ethernet segments at both ends is transparent to shortest path tree (SPT) bridge protocol data units (BPDUs), Virtual Terminal Protocol (VTP) packets, and other control messages.

The attachment circuit can be the Ethernet port or 802.1Q subinterface (VLAN). For each attachment circuit, LDP signals a different VC type via the targeted LDP session. VC type 5 is used for Ethernet port mode and VC type 4 is used for Ethernet VLAN mode.

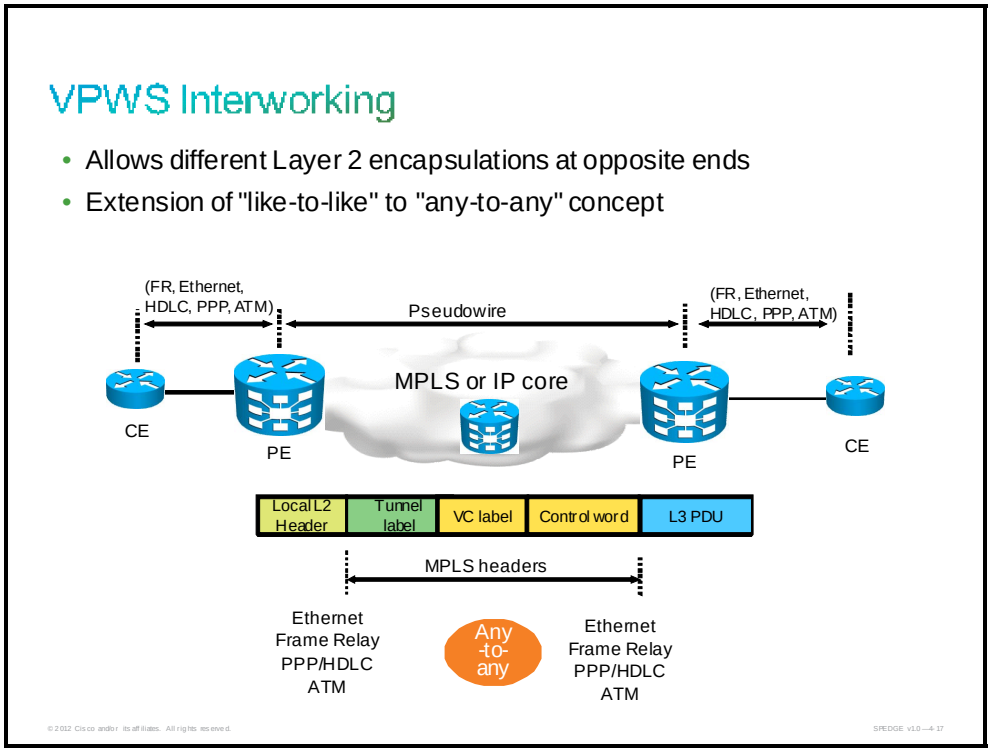
In Ethernet port mode, both ends of PW are connected to Ethernet ports. In this mode, the port is tunneled over PW or, using local switching (also known as attachment circuit-to-attachment circuit cross-connect) switches packets or frames from one attachment circuit to another attached to the same PE node. In Ethernet port mode, the PW is always a type 5 virtual connection. On the ingress PE, the network service provider passes the packets to the PW termination point, adds the MPLS labels to the packets, and sends the packets over the PW. In Ethernet port mode, a VLAN header may or may not be present in the frame. In any case, the PE router carries the frame transparently. This allows an Ethernet trunk to be carried over a single PW.

VLAN mode provides Ethernet VLAN-to-VLAN connectivity. In VLAN mode, each VLAN on a customer-end to provider-end link can be configured as a separate Layer 2 VPN connection, using either virtual connection type 4 or type 5.

Virtual connection type 5 is the default mode. In type 4 virtual connections, on the ingress PE, the VLAN tag maps to a particular PW and the packet is placed on the PW with the VLAN tag untouched.

In Type 5 virtual connections, on the ingress PE that is receiving packets from the customer edge (CE), the network service provider strips off the customer edge VLAN tag before placing the packets on the PW. On the egress PE, the network service provider pushes the VLAN tag onto the protocol stack before it sends the packet to the CE.

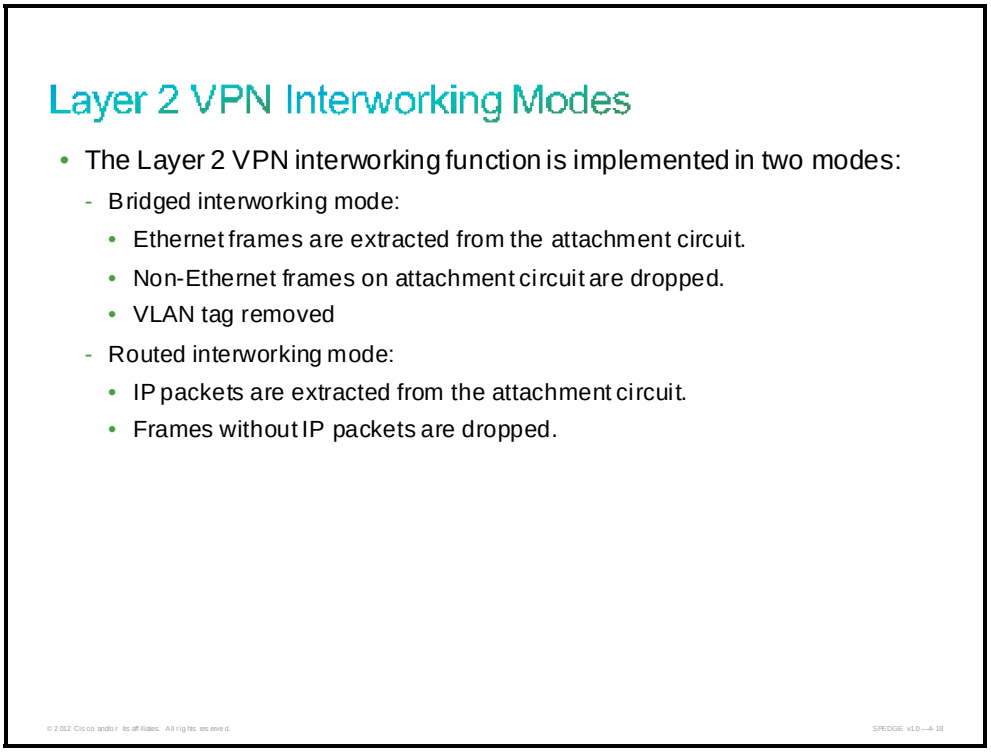
https://t.me/learningnets



Layer 2 VPN interworking allows you to connect disparate attachment circuits. Cisco routers support these any-to-any combinations:

- Ethernet or VLAN to ATM AAL5 interworking
- Ethernet or VLAN to Frame Relay interworking
- Ethernet or VLAN to PPP interworking
- Ethernet to VLAN interworking
- Frame Relay to ATM AAL5 interworking
- Frame Relay to PPP interworking
- Ethernet or VLAN to ATM VPI and VCI interworking

https://t.me/learningnets



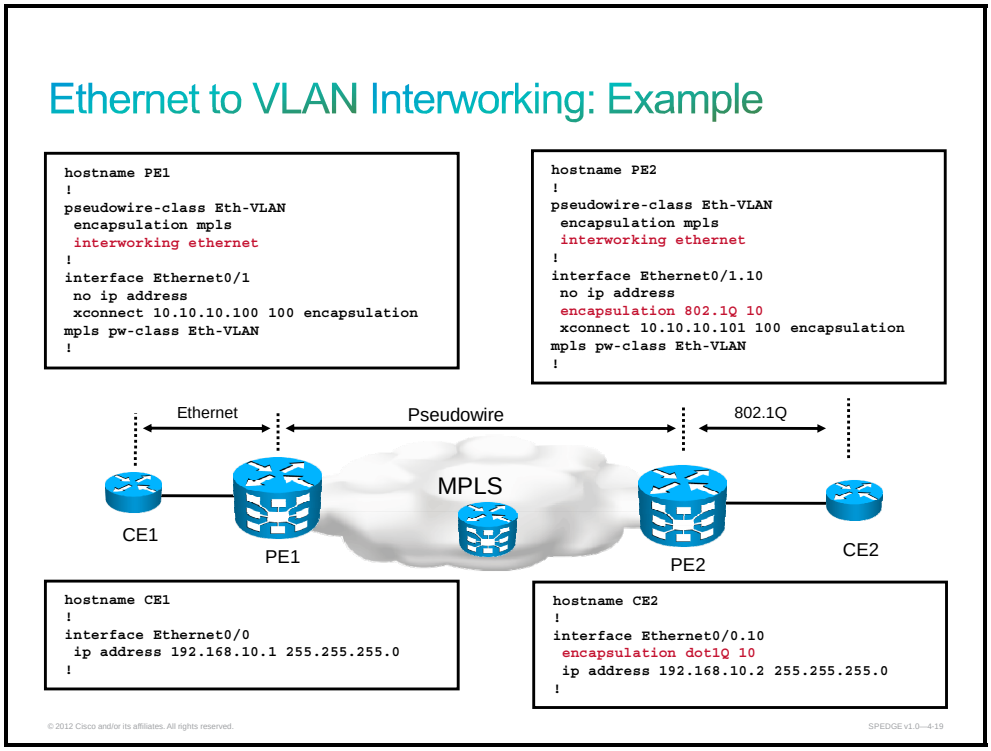
The L2VPN interworking function is implemented in two modes:

- Bridged interworking mode
- Routed interworking mode

In bridged interworking mode, Ethernet frames are extracted from the attachment circuit and sent over the PW. Attachment circuit frames that are not Ethernet are dropped. In the case of a VLAN, the VLAN tag is removed, leaving an untagged Ethernet frame. This interworking functionality is implemented by configuring the **interworking ethernet** command under the pseudowire class configuration mode.

In routed interworking, IP packets are extracted from the attachment circuit and sent over the PW. Attachment circuit frames are dropped if they do not contain the IPv4 packets. This interworking functionality is implemented by configuring the **interworking ip** command under the pseudowire class configuration mode.

https://t.me/learningnets



The steps to configure Ethernet to VLAN interworking between CE1 and CE2 are as follows:

Step 1 **Define pseudowire class on PE routers**

In this step, a pseudowire class called Eth-VLAN is defined on the PE1 and PE2 routers. This class configures the PW between the PE routers PE1 and PE2. Ensure that the parameters of the pseudowire class are the same on both PEs to enable PW establishment. The example shows that AToM encapsulation (encapsulation mpls) and bridged interworking mode (interworking Ethernet) will be used by the pseudowire class on the PE routers PE1 and PE2.

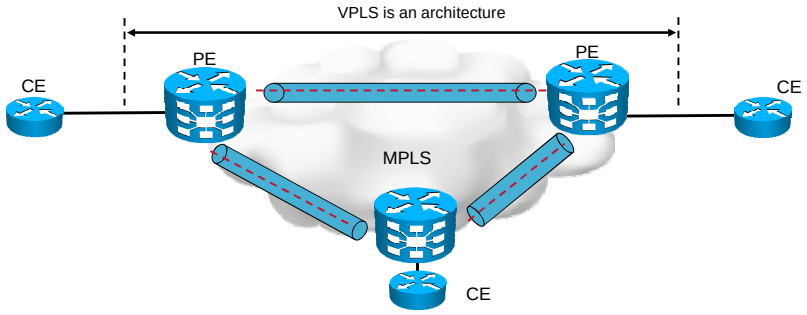
Step 2 **Define AToM VC to transport Layer 2 frames**

In this step, use the **xconnect** statement to define the AToM VC to carry the Layer 2 frames from CE1 to CE2, and vice versa. Associate the pseudowire class defined in Step 1 with the AToM VC.

https://t.me/learningnets

Virtual Private LAN Service

- End-to-end architecture
- Layer 2 multipoint Ethernet service:
  - MPLS transport (not L2TPv3)
  - Virtual bridges linked with PWs
- Service provider emulates an IEEE Ethernet bridge network
- Same data plane as EoMPLS (point-to-point)



© 2012 Cisco and/or its affiliates. All rights reserved.

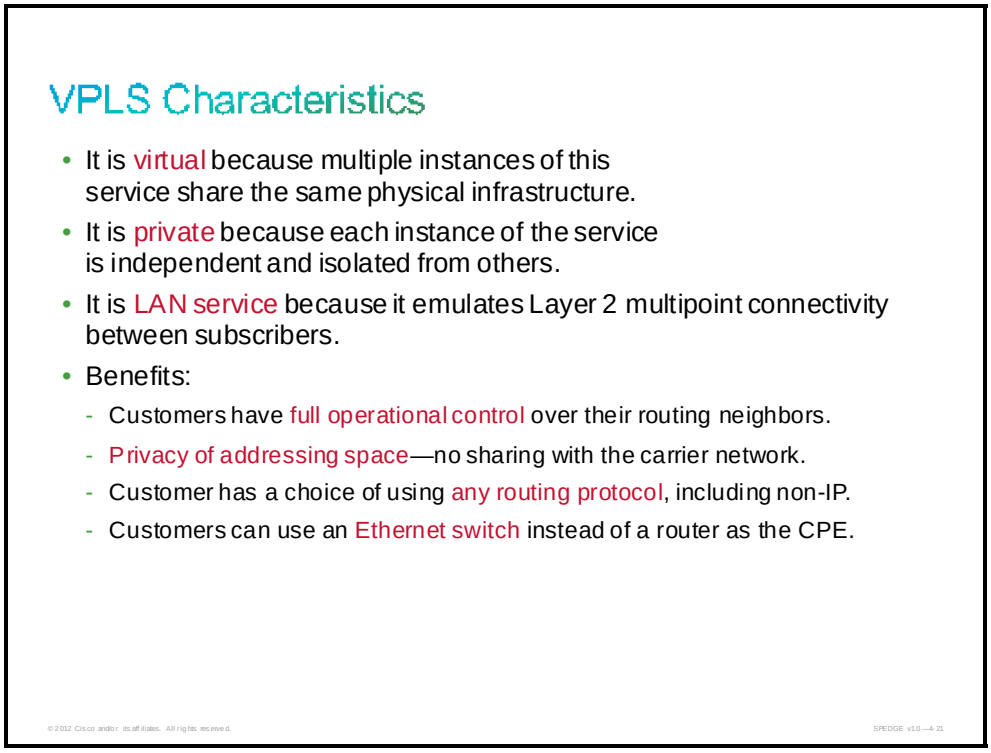
The purpose of VPLS is to provide a private multipoint LAN-type Ethernet connectivity service. VPLS emulates a LAN segment over an MPLS backbone across PWs or virtual circuits (VCs). VPLS creates one or more LANs for each customer who is using the service from the service provider. Each LAN is completely separate from the other emulated LAN segments. When a customer with different Ethernet sites connects to an MPLS backbone where VPLS is deployed, it appears as if all the sites are interconnected through a virtual Ethernet switch.

For each VPLS, the PE routers are fully meshed with PWs. A PE receiving a frame from another PE can identify which VPLS the frame belongs to, based on a PW label or VC label. As far as each customer is concerned, an Ethernet frame that is sent into the service provider network is delivered to the correct site(s) based on the destination MAC address. It is the task of each PE router to inspect the destination MAC address of each frame arriving from a locally attached site and to forward it to the appropriate destination site. This destination site may be attached to the same PE on a different port or a remote PE. If the destination site is attached to the same PE, the PE locally switches the frame to the correct port. If the destination site is attached to a remote PE, the ingress PE must forward the frame to the appropriate PW to the remote PE. This means that the ingress PE needs to know which egress PE to send the frame to.

There are two ways in which this can be accomplished. One is to have a control plane signaling to carry information about MAC addresses between PEs; another is to have a scheme based on MAC address learning. VPLS takes the latter approach by having each PE take the responsibility for learning which remote PE is associated with a given MAC address. This way, an ingress PE simply needs to identify which frames need to be sent to egress PEs, and egress PEs take care of identifying which local ports to forward the packet to. By inspecting the source MAC address of the frame arriving on a port, whether an actual local port or a PW from a remote PE, and by creating a corresponding entry in the forwarding table, the PE learns where to send future frames with that destination MAC address.

If Ethernet switches are used as CE devices and connected to PE routers, the PEs need to learn the MAC addresses of individual hosts attached to the switches. So, if a host is plugged into the office network served by a switch as a CE, the effect will be felt by all PEs. Thus, for a large deployment, it is better to use routers as CEs than switches.

https://t.me/learningnets



VPLS has become a very attractive technology over the past few years with the advent of MPLS. The reason for this is that some enterprises are very reluctant to relinquish the routing control of their network to the service provider, and they desire Layer 2 VPN services with multipoint connectivity. VPLS allows service providers to deploy carrier-class service over Ethernet and MPLS-based networks in a reliable and flexible way.

The term implies these characteristics:

- It is “virtual” because multiple instances of this service share the same physical infrastructure.
- It is “private” because each instance of the service is independent from the others.
- It is “LAN service” because it emulates Layer 2 multipoint connectivity between subscribers.

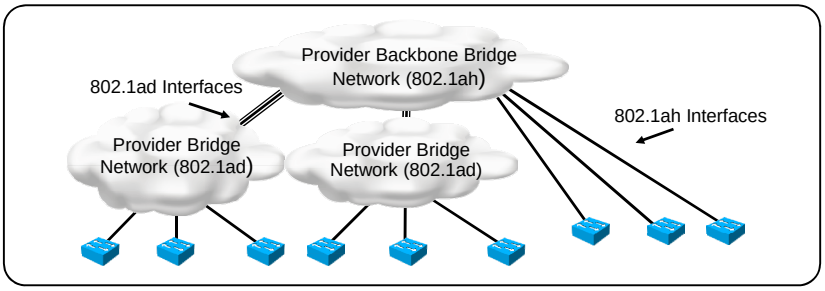
The main VPLS benefits include the following:

- Customers have full operational control over their routing neighbors.
- Privacy of addressing space means that addresses of the carrier infrastructure are completely isolated.
- Customers have a choice of using any routing protocol, including non-IP protocols such as Intermediate System-to-Intermediate System (IS-IS).
- Customers can use an Ethernet switch instead of a router as the customer premises equipment (CPE) device.

https://t.me/learningnets

802.1ah (Provider Backbone Bridges)

- Scales multipoint Layer 2 services:
  - 16 million service IDs
- Customer demarcation
- MAC hiding
- Flooding elimination
- VPN aggregation



© 2012 Cisco and/or its affiliates. All rights reserved.

Because of the lack of separation of customer networks from the carrier network (control plane), lack of scalability (limit of 4094 VLAN IDs), and lack of end-to-end QoS needed to achieve connection-oriented, carrier-grade Ethernet services, the standard known as provider backbone bridges (PBBs), or 802.1ah, was developed.

802.1ad provides stacking of VLAN IDs and allow separation of the customer VLAN ID from the service provider VLAN ID.

But since the control plane operates at the MAC layer and the goal is to provide separation of the customer control plane from the service provider control plane, one way is to simply “stack” the MAC addresses in a similar manner. This “stack” MAC approach is defined in the 802.1ah PBB standard, which is also referred to as MAC in MAC.

PBB is a set of architecture and protocols for routing over a provider network, allowing the interconnection of multiple provider bridge networks without losing the individually defined VLANs of each customer. It provides a further enhancement over QinQ tunneling to support even larger Ethernet deployments. QinQ does not offer true separation of customer and provider domains, but is merely a way to overcome the limitations on the VLAN identifier space.

The idea of PBB is to offer complete separation of customer and provider domains. It addresses the constraints of 802.1ad, such as having too little control on the MAC addresses, since QinQ forwarding is still based on the customer destination addresses. PBB eliminates flooding from the provider infrastructure and allows an efficient VPN aggregation.

802.1ad and 802.1ah can still be used hand-in-hand, as shown in the figure. QinQ is commonly used in the edge network, while PBB is deployed in the core.

PBB defines a new Ethernet header. The main components of the header are as follows:

- Backbone component that has:
  - Backbone destination address (B-DA) (6 bytes)
  - Backbone source address (B-SA) (6 bytes)

https://t.me/learningnets

- EtherType 0x88A8 (2 bytes)
- Backbone VLAN tag (B-TAG) and backbone VLAN ID (B-VID) (2 bytes); this is the backbone VLAN indicator

- Service encapsulation that has:
  - EtherType 0x88E7 (2 bytes)
  - Flags that contain priority, drop eligible indicator (DEI), and No Customer Address indication (for example, OAM frames).
  - Service instance VLAN ID (I-SID) (3 bytes)
- Original customer frame
  - Customer source address (6 bytes)
  - Customer destination address (6 bytes)
  - EtherType 0x8100 (2 bytes)
  - Customer VLAN identifier (2 bytes)
  - EtherType (e.g. 0x0800)
  - Customer payload

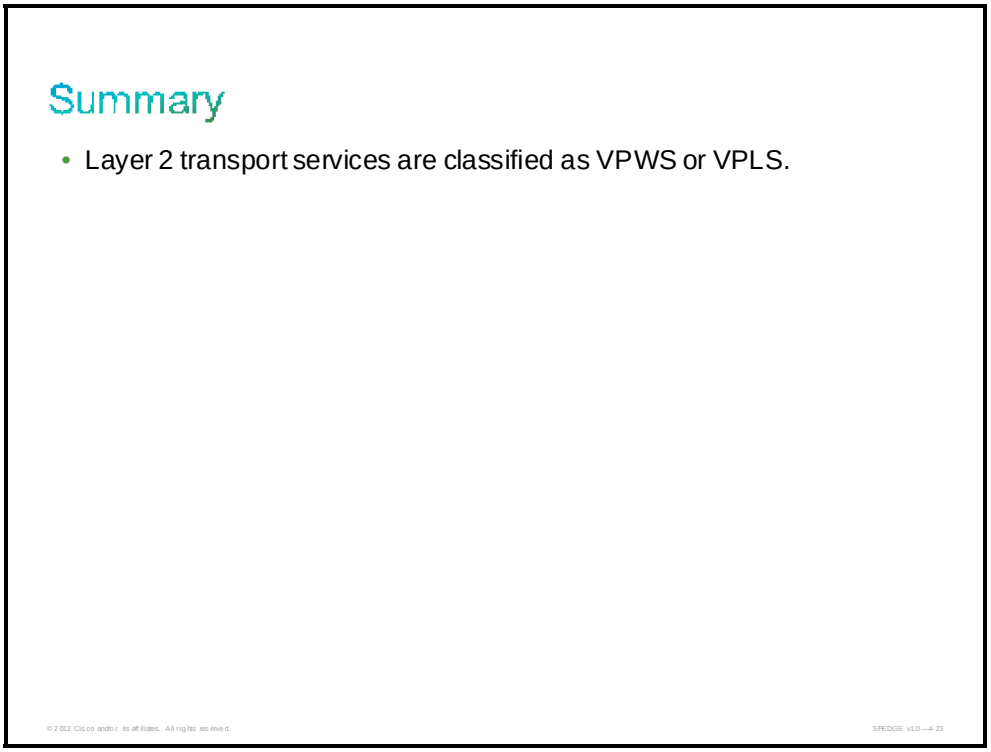
PBB defines a 48-bit B-DA and 48-bit B-SA to indicate the backbone source and destination MAC addresses. It also defines a 12-bit B-VID and 24-bit I-SID. The bridges in the PBB domain switch based on the B-VID and B-DA values, which contain 60 bits total. Bridges learn based on the B-SA and ingress port value and hence are completely unaware of the customer MAC addresses. I-SID allows for distinguishing the services within a PBB domain.

Note Detailed information on PBB implementations is beyond the scope of this class.

https://t.me/learningnets

Summary

This topic summarizes the primary point that was discussed in this lesson.



https://t.me/learningnets

Lesson 2

Introducing AToM

Overview

Any Transport over MPLS (AToM) is a subset of Virtual Private Wire Service (VPWS) that provides point-to-point virtual connections. AToM allows an MPLS network to provide end-to-end transport for Layer 2 frames and cells. It provides support for Ethernet, PPP, High-Level Data Link Control (HDLC), Frame Relay, and ATM. This lesson describes AToM characteristics, implementation, and verification.

Objectives

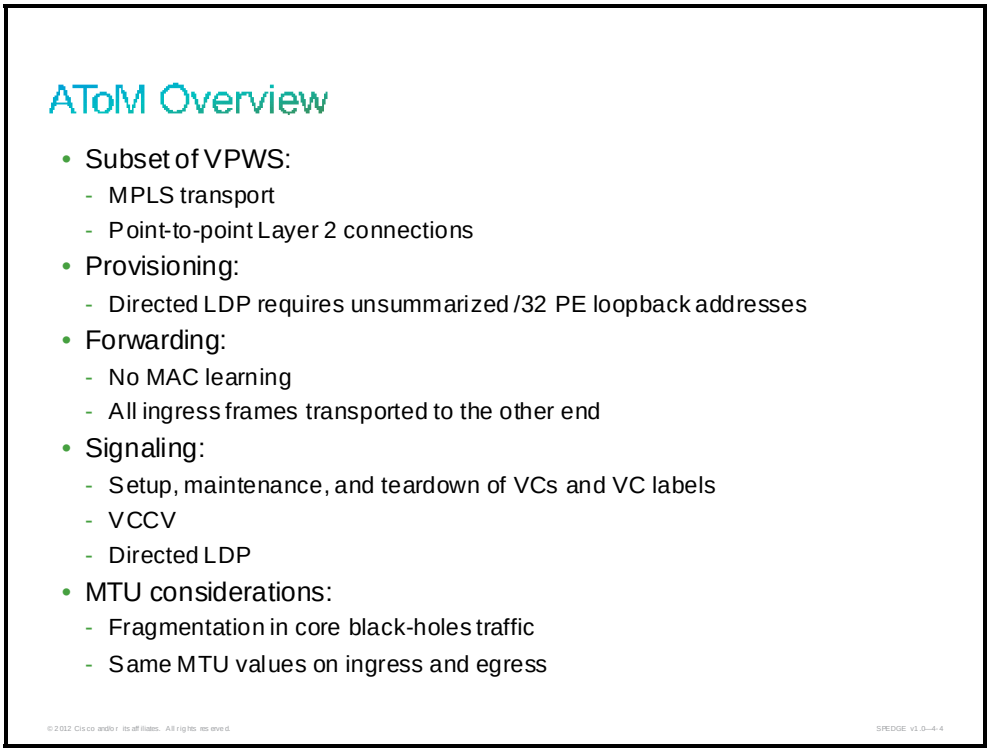
Upon completing this lesson, you will be able to describe AToM. You will be able to meet these objectives:

- Introduce AToM
- Implement AToM

https://t.me/learningnets

Introduction to AToM

This topic explains AToM.



AToM is a subset of Virtual VPWS that enables point-to-point Layer 2 virtual connections over an MPLS infrastructure. Several aspects require special consideration:

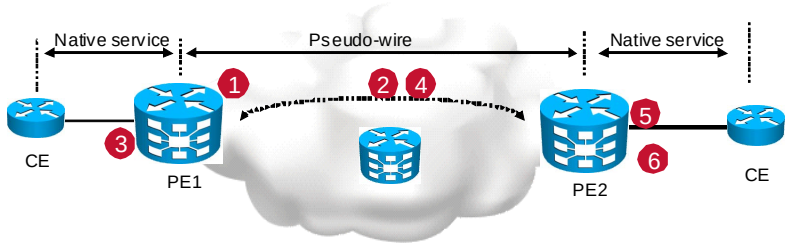
- Provisioning
- Forwarding
- Signaling
- Maximum transmission unit (MTU) considerations

These aspects are discussed in this topic.

https://t.me/learningnets

AToM Provisioning

1. Use the **xconnect** command on ingress PE (port, subinterface, etc.).
2. PE1 starts a directed LDP session to PE2 (if not yet available):
   - One LDP session can signal multiple PWs.
3. PE1 allocates the VC label and binds to the VC ID:
   - Same VC ID on both ends; VC label unique per PE
4. PE1 sends mapping message (VC FEC TLV, VC label TLV).
5. PE2 receives VC FEC and label TLV and maps to local VC ID.
6. PE2 repeats the process (1 to 4, and then 5 on PE1).



© 2012 Cisco and/or its affiliates. All rights reserved.

AToM provisioning occurs in these steps:

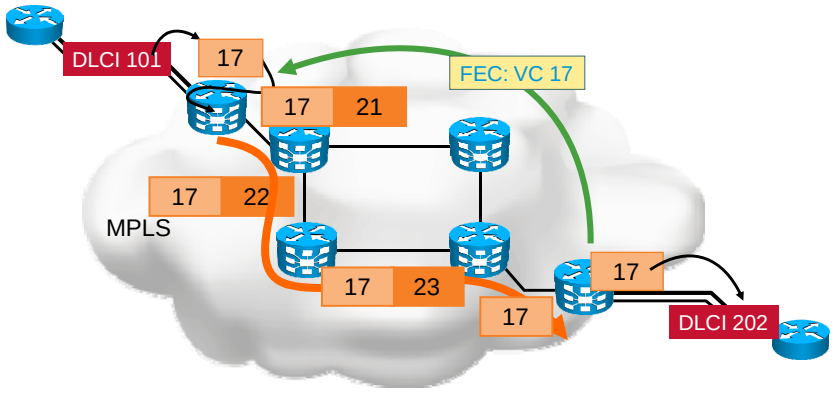
1. The **xconnect** command is issued on the ingress provider edge (PE) (port, subinterface, and so on). Alternatively, the pseudowire (PW) autodiscovery mechanism can be used to detect the neighbor.
2. PE1 starts a directed Label Distribution Protocol (LDP) session to PE2 if one is not yet available. If a directed session already exists to the destination PE, it is reused for another PW. One LDP session can signal multiple PWs.
3. PE1 allocates a virtual circuit (VC) label and binds it to the VC ID. The same VC ID value must be configured on both ends. The VC label is unique per PE.
4. PE1 sends a mapping message containing the VC forwarding equivalent class (FEC) type, length, and value (TLV) and VC label TLV to the other end.
5. The other end (PE2) receives VC FEC and label TLV and maps it to the locally configured VC ID.
6. The other end (PE2) repeats the process (Steps 1-4), which is then finished on PE1 by receiving the VC FEC and label TLV and mapping it to the locally configured VC ID.

https://t.me/learningnets

AToM Forwarding

FEC

- Set of packets handled in the same way on MPLS LSR
- Used to bind a VC label to a VC ID
- Multiplexing customer data over the same LSP tunnel



© 2012 Cisco and/or its affiliates. All rights reserved.

An Interior Gateway Protocol (IGP) is required in the MPLS backbone. All routers in the backbone have routing information about how to send IP packets to each other. LDP is also used between directly connected neighbors. Local labels are assigned to each IGP-derived route. The label values are then propagated to the neighbor across the LDP session.

The IGP, together with the LDP sessions between directly connected neighbors, establishes label-switched paths (LSPs) from any router inside the backbone to any other router inside the backbone. In the figure, a unidirectional LSP from the upper right to the lower left is established between the ingress and egress PE routers. The tunnel label is used to propagate the packets along the LSP to the correct egress PE router.

The figure also shows the directed multihop LDP session that is used to exchange the VC label between the ingress and egress PE routers. Any ingress-egress PE router pair will need such an LDP session. In this example, the egress PE router allocates the label value 17. The VC label is advertised to the ingress PE router using the directed LDP session between them.

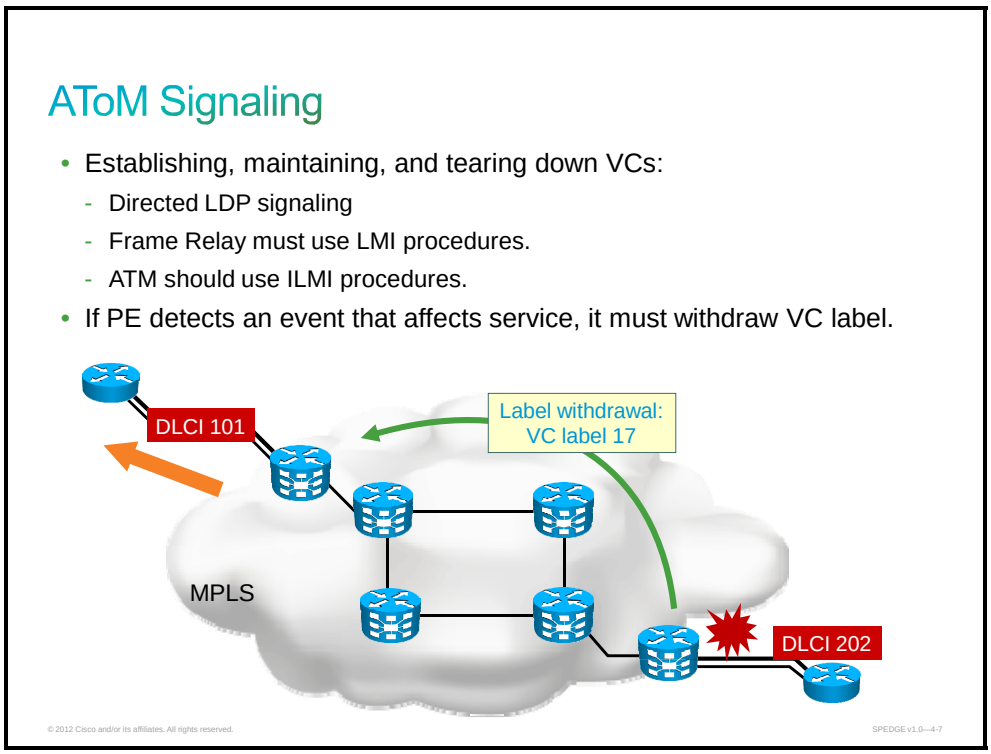
The ingress PE router now forms a label stack. The topmost label, the tunnel label, has the value 21 and is used to guide the packets to the egress PE router. The second label, the VC label, has the value 17 and is used by the egress PE router to propagate out the packets on the correct interface.

The ingress PE router receives a Frame Relay frame on data-link connection identifier (DLCI) 101 on the incoming interface. The DLCI is mapped to the AToM tunnel across the backbone. The Frame Relay frame is therefore encapsulated into MPLS using the label stack with label 21 as the topmost label and label 17 as the second label.

The packet is then forwarded along the LSP. The topmost label is used for label swapping in the next hop. The top label is changed to the value 22. In the next hop, label swapping results in label value 23 being the top label. In the router just before the egress router, the incoming label value 23 indicates pop. That label therefore performs penultimate hop popping (PHP). The topmost label is removed, and the packet is propagated to the egress PE router with the label value 17, the VC label, which is now the only label left.

https://t.me/learningnets

When the PE router receives the packet with label value 17, that label value instructs the PE router to decapsulate the packet and send it out on the associated Frame Relay DLCI. In this case, the DLCI value is 202. The Frame Relay frame is now reconstructed and transmitted.



The signaling of the VC occurs over the targeted LDP sessions and includes the setup, maintenance, and teardown of the PW. In addition, the PE routers translate the events from the local attachment circuit to the PW and vice versa. This translation depends on the encapsulation of the local attachment circuit.

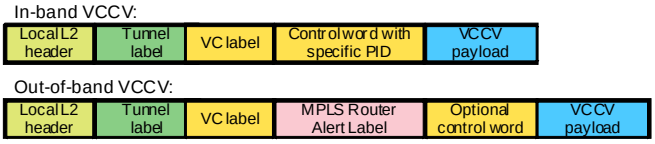
A PE router may provide circuit status signaling on the interface where the customer connects. A PE router that provides Frame Relay services must use Local Management Interface (LMI) procedures with the customer equipment. A PE router that provides ATM services should use Integrated Local Management Interface (ILMI) procedures.

An example of the signaling procedure, shown here, is the VC teardown. If a PE router detects a condition that affects normal service, it must withdraw the corresponding VC label.

https://t.me/learningnets

Virtual Circuit Connectivity Verification

- Layer 2 VPN OAM feature
- Keepalive protocol to monitor PW data forwarding
- AToM VCCV categories:
  - Switching modes—for differentiating between control and data traffic
    - In-band (type 1) —uses PID field in the AToM control word to identify VCCV control packet
    - Out-of-band (type 2) —MPLS router alert label is carried above the VC label to identify VCCV control packet
  - Applications—in-band keepalive method
    - MPLS LSP ping
    - ICMP ping



© 2012 Cisco and/or its affiliates. All rights reserved.

Virtual Circuit Connectivity Verification (VCCV) is Layer 2 VPN Operation, Administration, and Maintenance (OAM) feature that allows network operators to run a PE-to-PE keepalive protocol across a specified PW to ensure that the PW data path forwarding does not contain any faults.

When a PW is first signaled using LDP or L2TPv3, a message is sent from the initiating PE to the receiving PE. That message has been extended to include VCCV capability information that indicates to the receiving PE which combinations of control channel and connectivity verification types it is capable of receiving. If the receiving PE agrees to establish the PW, it will return its capabilities in the subsequent signaling message.

When MPLS is used to transport PW packets, VCCV packets are carried over the MPLS LSP. Packets are sent across this channel either as in-band traffic with the data of the PW, or out-of-band.

Two types of VCCV switching modes have been defined to distinguish VCCV packets from regular data traffic:

- **In-band VCCV (type 1):** Control channel type 1 uses a protocol ID (PID) field in the AToM control word to identify an AToM VCCV packet. This type of VCCV is used for those PW types that employ the control word.
- **Out-of-band VCCV (type 2):** Control channel type 2 is also referred to as the MPLS router alert label. The MPLS router alert label is carried above the VC label to identify an AToM VCCV packet. Control channel type 2 can be used whether the PW is set up with a control word present or not.

Cisco routers use type 1 switching, if available, when they send MPLS LSP ping packets over an AToM VC control channel. Type 2 switching accommodates those VC types and implementations that do not support or interpret the AToM control word.

https://t.me/learningnets

An AToM VC advertises its AToM VCCV disposition capabilities in both directions—that is, from the originating router (PE1) to the destination router (PE2), and from PE2 to PE1. In some instances, AToM VCs might use different switching types if the two endpoints have different AToM VCCV capabilities. If PE1 supports type 1 and type 2 AToM VCCV switching and PE2 supports only type 2 AToM VCCV switching, there are two consequences:

- LSP ping packets sent from PE1 to PE2 are encapsulated with type 2 switching.
- LSP ping packets sent from PE2 to PE1 use type 1 switching.

You can determine the AToM VCCV capabilities advertised to and received from the peer by entering the **show mpls l2transport binding** command (Cisco IOS Software) or **show l2vpn xconnect all detail** command (Cisco IOS XR Software).

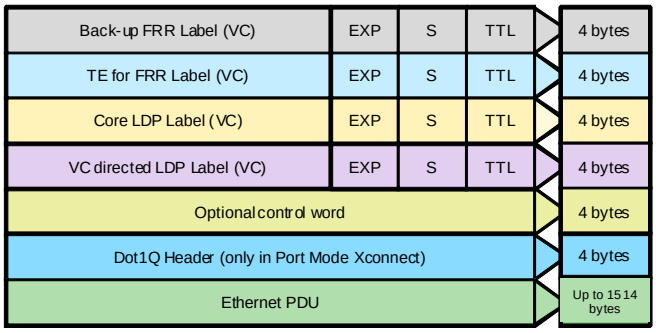
Two verification types can be negotiated by the in-band or out-of-band VCCV control channel in an MPLS environment:

- **ICMP ping:** When this optional connectivity verification mode is used, an Internet Control Message Protocol (ICMP) echo packet (ICMPv4 or ICMPv6) achieves connectivity verification.
- **MPLS LSP ping:** This method helps monitor LSPs and quickly isolate MPLS forwarding problems.

https://t.me/learningnets

AToM MTU Considerations

- AToM transport of Frame Relay, Ethernet, and AAL5 does not allow packets to be fragmented and reassembled.
- Ensure that the MTU of all intermediate links between endpoints is sufficient to carry the largest Layer 2 frame received.
- The ingress and egress PE routers must have the same MTU value.



© 2012 Cisco and/or its affiliates. All rights reserved.

Unlike IP, most Layer 2 protocols (for example, Frame Relay, Ethernet, and ATM adaptation layer 5 [AAL5]) do not allow fragmentation of frames. This fact has two implications:

- All intermediate links between the ingress PE router and the egress PE router must be able to carry the largest Layer 2 frame that has been received, including the imposed label stack and the 4-byte control word (if it is used).
- The ingress PE interface and the egress PE interface must have the same MTU value.

Failure to comply with the first rule means that the larger frames, where the label stack and the control word contribute to creating a larger size than can be carried, will be dropped by the backbone.

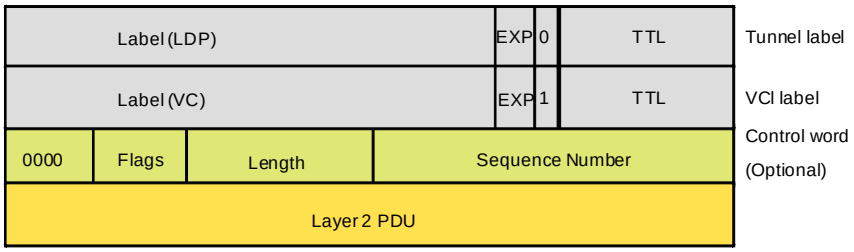
Failure to comply with the second rule means that frames that are forwarded along the LSP will be dropped by the egress PE if the frame size is too large for the egress PE interface.

The figure illustrates the potential overhead imposed by various labels and headers. The MTU size on the core must be able to accommodate the customer MTU size with the added overhead.

https://t.me/learningnets

AToM Control Word

- Control word is optional.
- Transmitted after the label or labels and before the Layer 2 PDU
- Can be used for in-band VCCV
- Flag field carries different bits for different Layer 2 protocols:
  - Frame Relay: FECN, BECN, DE, C/R
  - ATM: AAL5 or cell, EFCI, CLP, C/R
- Sequence number 0 indicates that no sequencing is done.



© 2012 Cisco and/or its affiliates. All rights reserved.

The figure illustrates AToM encapsulation. The topmost label is the tunnel label. This label is followed by the VC label. Following the VC label is the optional control word. Next in the packet is the Layer 2 protocol data unit (PDU).

The control word is optional. Some Layer 2 protocols make use of it, while others do not. Both endpoints (ingress PE and egress PE) must agree to either use or not use the control word. It is transmitted after the label stack but before the Layer 2 PDU. It can be used to carry important Layer 2 header information and to guarantee sequenced delivery, if required.

The control word is 32 bits long. It is divided into four fields. The first field is 4 reserved bits, which must always be set to zero. The next field is a 4-bit flag field. The flags have different uses depending on the Layer 2 protocol that is being forwarded. The third field is an 8-bit length field, which is used only if the Layer 2 PDU is shorter than the minimum MPLS packet and padding is required. If no padding is required, the length field is not used. The fourth field is a 16-bit sequence number. Sequence numbering is used only on Layer 2 protocols, and it guarantees ordered delivery. A special value of 0 in the sequence field indicates that there is no guaranteed sequenced delivery.

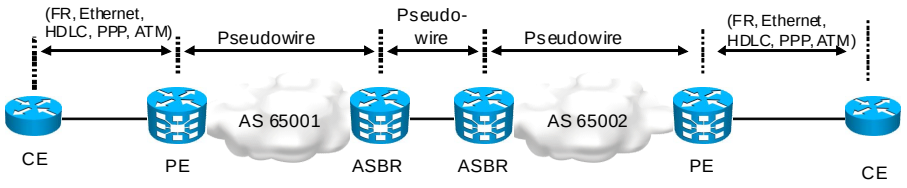
When AToM is used for Frame Relay over MPLS (FRoMPLS), the Frame Relay header is removed and the forward explicit congestion notification (FECN), backward explicit congestion notification (BECN), discard eligible (DE), and command/response (C/R) bits are carried in the control word flag field.

When AToM is used for ATM over MPLS, the first flag in the control word flag field is used to indicate whether it is AAL5 frames or raw ATM cells that are being transported by AToM. The other three flags are used for explicit forward congestion indication (EFCI), cell loss priority (CLP), and C/R. If one of these flags is set in any of the ATM cells that are being transported in the MPLS packet, then the corresponding flag is set in the control word.

https://t.me/learningnets

Inter-AS AToM

- Layer 2 VPN extends VCs over single service provider AS.
- Changes in control and data plane code are required for inter-AS span.
- PW stitching solution:
  - Interconnects PWs in different autonomous systems
  - ASBRs are the stitch points
  - Interworking of control and data planes at stitch point



© 2012 Cisco and/or its affiliates. All rights reserved.

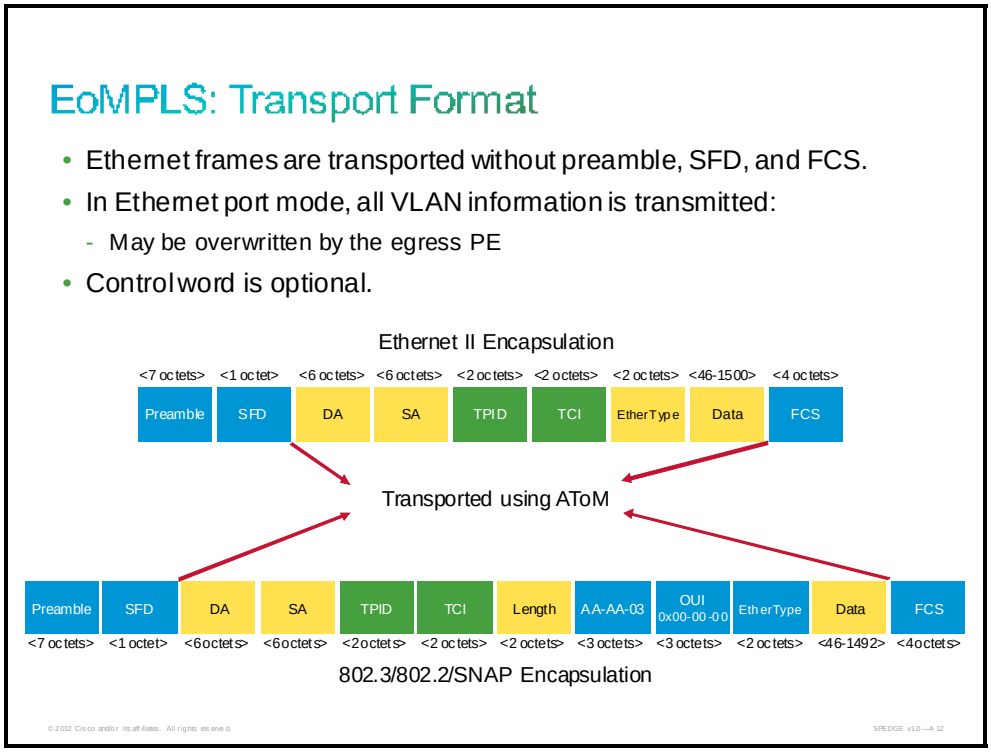
Traditionally, AToM enables VC connections through a single autonomous system (AS). To extend this function to interautonomous system (inter-AS) deployments, changes in the control and data would be required. Pseudowire (PW) stitching provides an immediate solution.

PW stitching enables the extension of Layer 2 VPN PWs across an inter-AS boundary or across two separate MPLS networks. Layer 2 VPN PW stitching connects two or more contiguous PW segments to form an end-to-end multihop PW. This end-to-end PW functions as a single point-to-point PW. Layer 2 VPN PW stitching enables the service provider to keep the IP addresses of the edge PE routers private across inter-AS boundaries. Using the IP addresses of the Autonomous System Boundary Routers (ASBRs), the ASBRs join the PWs of the two domains.

AToM packets forwarded between two PWs are treated the same as any other MPLS packet, with the following exceptions:

- The outgoing VC label replaces the incoming VC label in the packet. New IGP labels and Layer 2 encapsulation are added.
- The incoming VC label Time to Live (TTL) field is decremented by one and copied to the outgoing VC label TTL field.
- The incoming VC label experimental (EXP) value is copied to the outgoing VC label EXP field.
- The outgoing VC label “bottom-of-stack” S bit in the outgoing VC label is set to 1.
- AToM control word processing is not performed at the Layer 2 VPN PW switching aggregation point or ASBR. Sequence numbers are not validated.

https://t.me/learningnets



AToM encapsulates entire frames but excludes fields whose transmission does not offer any benefits, such as frame synchronization data. In Ethernet over MPLS (EoMPLS), the preamble, the start frame delimiter (SFD) and the frame check sequence (FCS) are excluded from encapsulation.

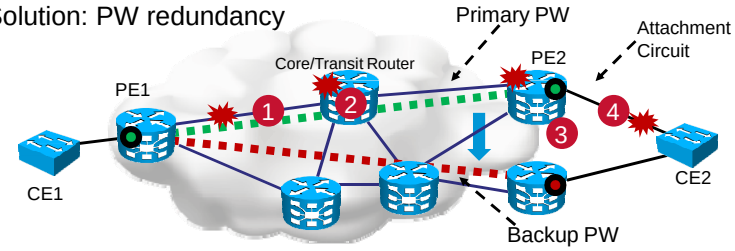
The preamble of an Ethernet frame consists of a 56-bit (7-byte) pattern of alternating 1 and 0 bits, which allows devices on the network to easily detect a new incoming frame. The SFD is designed to break this pattern, and signal the start of the actual frame. The SFD is the 8-bit (1-byte) value marking the end of the preamble of an Ethernet frame. The SFD is immediately followed by the destination MAC address. It has the value 10101011. The FCS provides Layer 2 checksum characters added to a frame for error detection and correction purposes.

In Ethernet port mode, all VLAN information is transmitted. The VLAN tag or tag stack may be overwritten by the egress PE. The egress PE can also manipulate the VLAN tag received over the VC. The control word is optional in EoMPLS.

https://t.me/learningnets

EoMPLS Redundancy

- Failures 1 and 2 (transit network):
  - IGP and MPLS LDP will reconverge.
  - With MPLS traffic engineering and FRR enabled, failover to backup tunnel.
  - PW will stay up as long as PE1 has available LSP path to PE2.
  - PW service layer is not affected.
- Failures 3 and 4 (service node or attachment circuit):
  - EoMPLS PW will go down.
  - Network transport layer reconverge does not help.
- Solution: PW redundancy



© 2012 Cisco and/or its affiliates. All rights reserved.

It is common practice to add redundancy to the interconnect between two customer sites to avoid split-subnet scenarios and service interruption. The redundancy solution comprises several building blocks. One block refers to the redundant LAN access and considers issues related to the Spanning Tree Protocol (STP). This component is described in the next lesson.

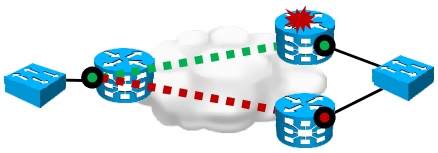
This figure presents a redundant transit network and redundant attachment circuits. In case of failure 1 or 2, as shown above, IGP and MPLS LDP will reconverge. MPLS traffic engineering and Fast Reroute (FRR) will trigger an automatic failover to the backup tunnel. The PW will stay up as long as PE1 has an available LSP path to PE2. This failure scenario does not affect the PW service layer.

In case of failure 3 or 4, as shown above, the PW will go down. Network layer reconvergence will not help re-establish the service. A solution has been developed to address this problem. It is referred to as PW redundancy.

https://t.me/learningnets

One-Way Redundancy

- Dual-homing of one local PE to:
  - Two remote PEs
  - Two different attachment circuits on the same remote PE
- Two PWs: Primary and backup provide redundancy for a single attachment circuit or node.
- Faults on the primary PW cause failover to backup PW.



Case 1: Service node and attachment circuit protection



Case 2: Attachment circuit protection

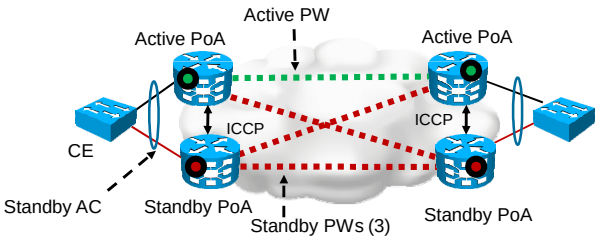
© 2012 Cisco and/or its affiliates. All rights reserved.

PW redundancy can be implemented using a one-way or two-way method. In one-way redundancy, the local PE has two PWs to two remote PEs serving the same destination site or to the same egress PE with redundant attachment circuits. One PW is declared as primary, the other as backup. A fault of the primary PW triggers a failover to the backup PW.

https://t.me/learningnets

Two-Way Redundancy

- Dual-homing of two local PEs to two remote PEs
- Four PWs:
  - One primary PW
  - Three backup PWs
- Requires MC-LAG
  - Point of attachment nodes run ICCP
  - ICCP synchronizes state and forms a redundancy group.



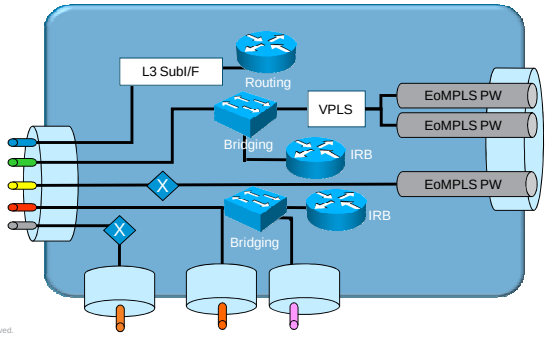
© 2012 Cisco and/or its affiliates. All rights reserved.

In two-way redundancy, four PWs are used to provide high availability service. Only one PW is declared as primary. The three remaining PWs are intended for backup. In each site, one attachment circuit is primary, the other backup. To synchronize the redundancy information between the LAN and the PE devices, Multi-Chassis Link Aggregation Group (MC-LAG) must be enabled in the network. MC-LAG refers to the PE devices as points of attachment nodes. The points of attachment run Inter-Chassis Communication Protocol (ICCP) to synchronize state and form a redundancy group.

https://t.me/learningnets

EVC Infrastructure

- Service instances configured on main interface:
  - Also known as EFP
- Each EFP matches a predefined VLAN tag-based criteria.
- Optional tag manipulation can be configured.
- Traffic forwarding is specified.
- Features such as QoS policies can be specified.



© 2012 Cisco and/or its affiliates. All rights reserved.

Ethernet Virtual Connection (EVC) is used to represent Cisco software architecture to address Carrier Ethernet services.

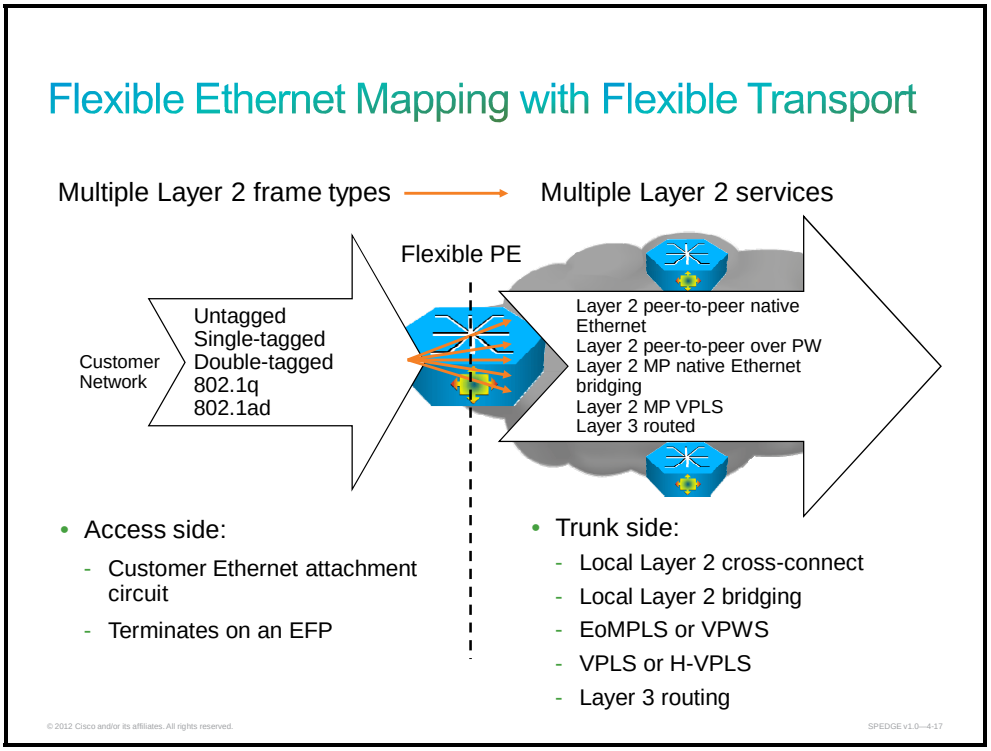
Cisco implements EVC using Ethernet Flow Points (EFPs). An EFP is a substream partition of a main interface. On Cisco routers, the EFP is implemented as a Layer 2 subinterface with an encapsulation statement.

The EVC solution defines these aspects of Ethernet-based attachment circuits and virtual circuits:

- Frame matching based on one or more VLAN tags
- Optional VLAN tag manipulation
- Traffic forwarding
- Additional services, such as quality of service (QoS) policies

EVC is not supported on all Cisco platforms. It is supported on Cisco IOS and IOS XR Software.

https://t.me/learningnets



Customer-to-service mapping is illustrated here as a connection between an Ethernet-based data flow on the customer side and a service on the trunk side.

Flexible Ethernet mapping is the ability to process and classify different Ethernet frame types, each with different attributes (EtherTypes, VLAN tags, class of service [CoS] bits, and so on). Cisco IOS XR Software uses the Ethernet Flow Point (EFP) concept to provide flexible Ethernet mapping. Flexible transport is found on the trunk side. Each Ethernet flow from the customer side is mapped or connected to a service on the trunk side. These service types can be native Ethernet, IP, or IP and MPLS-based, and they form the basis for Layer 2 VPN.

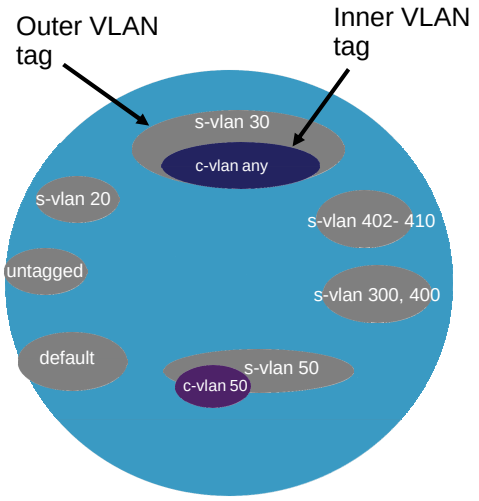
https://t.me/learningnets

Flexible Frame Matching and Service Mapping

EFPs enable flexible mapping of frames into Layer 2 services.

Mapping is based on VLAN tagging:

- 802.1Q, 802.1ad
- Single-tag or double-tag
- Unique or multiple values (ranges or lists)
- Untagged traffic
- Unclassified traffic (default)



© 2012 Cisco and/or its affiliates. All rights reserved.

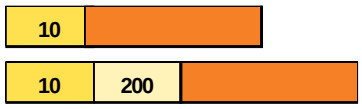
One of the EVC advantages is flexible frame matching. Flexible frame matching is a functionality that allows each service instance to match frames with either a unique single VLAN, or a list or range of VLANs. It can also match single- or double-tagged frames, untagged frames, or any frames that are not matched by the specific statement.

Flexible frame matching is the first step when configuring a service instance. Subsequent steps, encapsulation rewrite, and forwarding definition are described in the following pages.

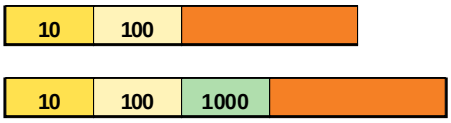
https://t.me/learningnets

EVC: Exact vs. Nonexact Match

- EVC supports only nonexact matching.
- **encapsulation dot1q 10** matches any packets with outmost tag equal to 10:



- **encapsulation dot1q 10 second 100** matches any packets with outmost tag equal to 10 and second outmost tag equal to 100:



© 2012 Cisco and/or its affiliates. All rights reserved.

Several subinterfaces with various matching statements can exist on a main interface. The matching logic affects the assignment of a frame to a specific subinterface. The EVC supports only nonexact matching, in which additional inner-VLAN tags are not taken into consideration.

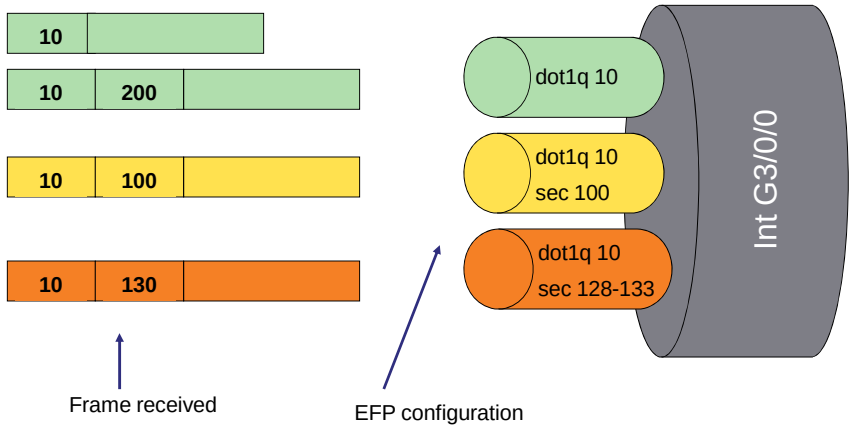
The command **encapsulation dot1q 10** matches any frames with the outmost VLAN tag equal to 10. This includes both 801.1q frames, queue-in-queue (QinQ) frames with any second VLAN tag, as well as frames with more than two VLAN tags, if the outmost tag is 10.

The command **encapsulation dot1q 10 second 100** matches any frames with the outmost VLAN tag equal to 10 and second outmost tag equal to 100. This includes frames that contain more than two VLAN tags, if the first two tags fulfill the defined criteria.

https://t.me/learningnets

EVC: Longest-Match Rule

- Longest match defines frame-to-EFP matching.



© 2012 Cisco and/or its affiliates. All rights reserved.

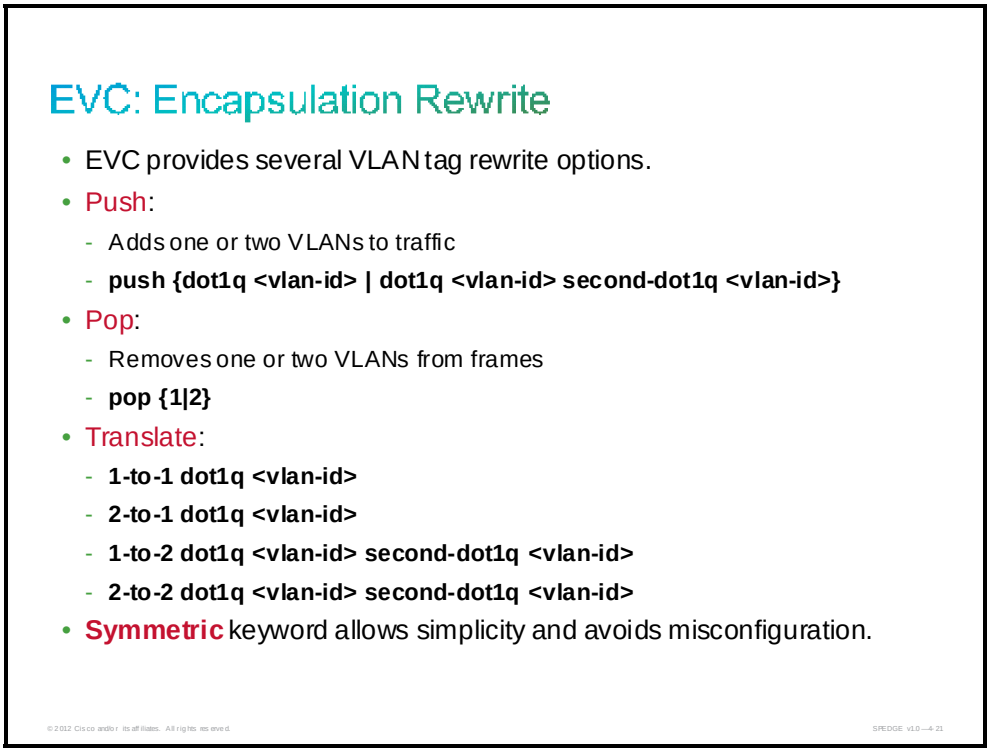
If several potential matches exist on the same main interfaces, the longest-match rule defines the most specific hit.

With the three subinterfaces shown above, the first one is the only subinterface that matches the first two frames (single tag 10, and double tag 10-200).

The second frame, with double tag (10-100) is matched by the first and second subinterface. Because the second subinterface defines a more exact match, it will process the second frame.

The third frame, with double tag (10-130) is matched by the first and third subinterface. Because the third subinterface defines a longer match, it will process the third frame.

https://t.me/learningnets

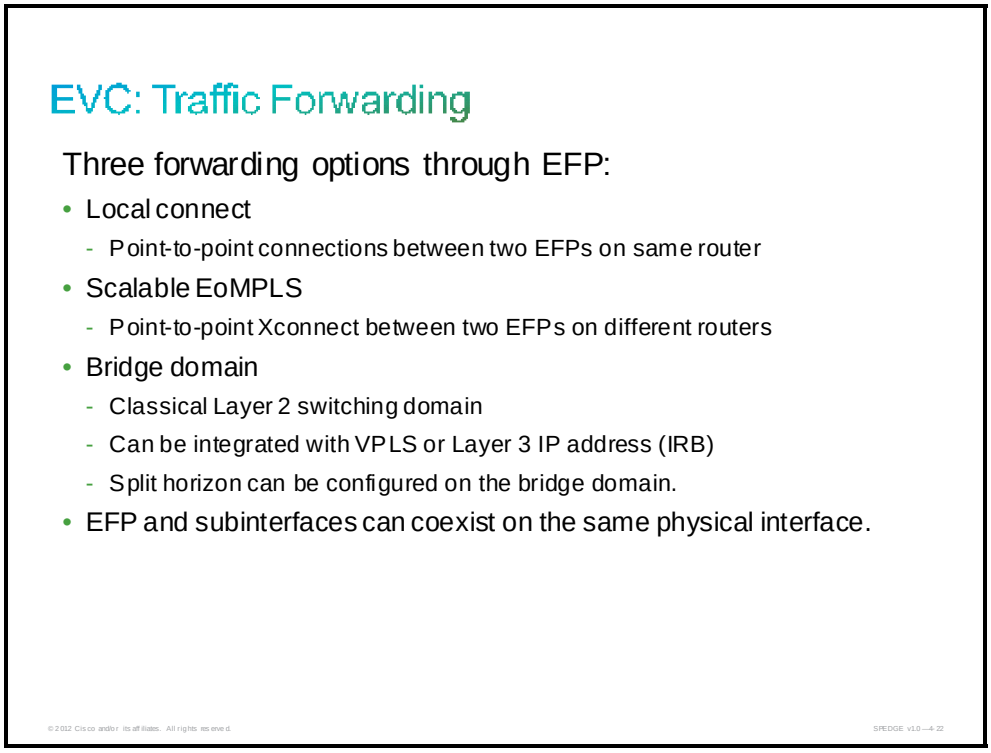


Once a packet is matched and assigned to a given subinterface, its VLAN tags can be modified. You can configure several of these three operations:

- **Push:** This action adds one or two VLANs to traffic using the **push {dot1q <vlan-id> | dot1q <vlan-id> second-dot1q <vlan-id>}** command.
- **Pop:** This action removes one or two VLANs from frames.
- **Translate:** This action allows the replacement of one or two tags by another one or two tags, using one of these commands:
  - **1-to-1 dot1q <vlan-id>**
  - **2-to-1 dot1q <vlan-id>**
  - **1-to-2 dot1q <vlan-id> second-dot1q <vlan-id>**
  - **2-to-2 dot1q <vlan-id> second-dot1q <vlan-id>**

The **symmetric** keyword defines that the pop, push, or translate operation is reversed for egress frames.

https://t.me/learningnets

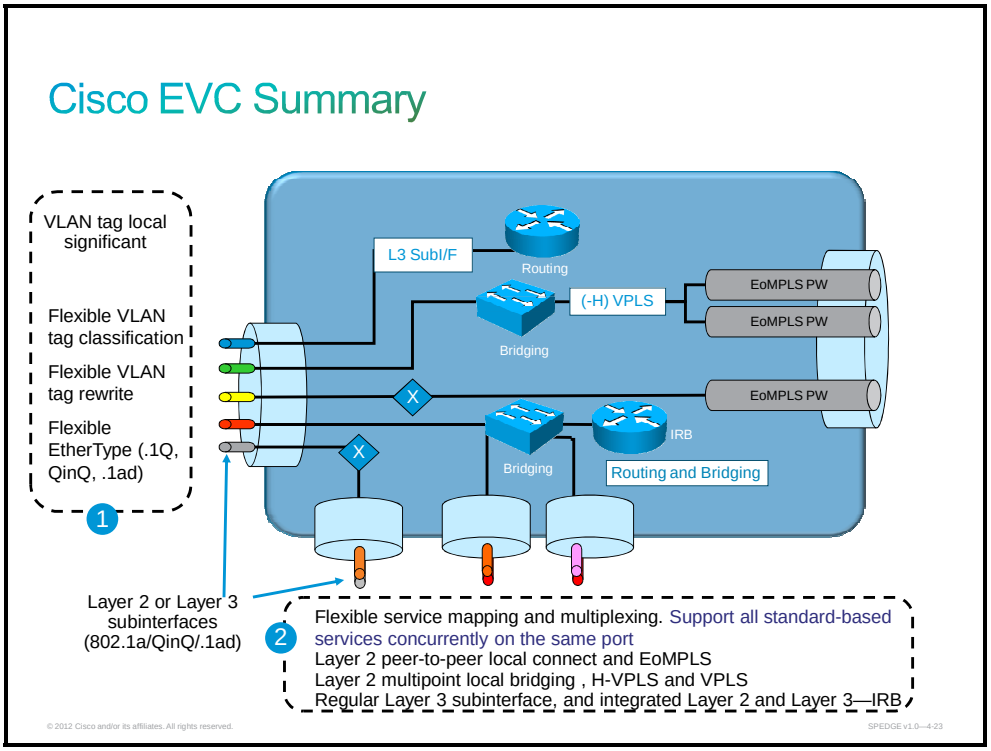


Lastly, the EVC framework is used to define traffic forwarding. Incoming frames can be transmitted to these destination types:

- **Local connect:** This destination represents a point-to-point connection between two EFPs on the same router. No VLANs are required to be defined on the PE device.
- **EoMPLS:** This destination refers to a point-to-point Xconnect between two EFPs on different PE routers. No VLANs are necessary on the PE device.
- **Bridge domain:** This destination describes a Virtual Private LAN Service (VPLS) that uses a classical Layer 2 switching domain. The associated switch virtual interface (SVI) can be enabled for a VPLS or Layer 3 IP address (integrated routing and bridging [IRB]). You can enable or disable the split-horizon rule within the bridge domain.

EFP and subinterfaces can coexist on the same physical interface.

https://t.me/learningnets

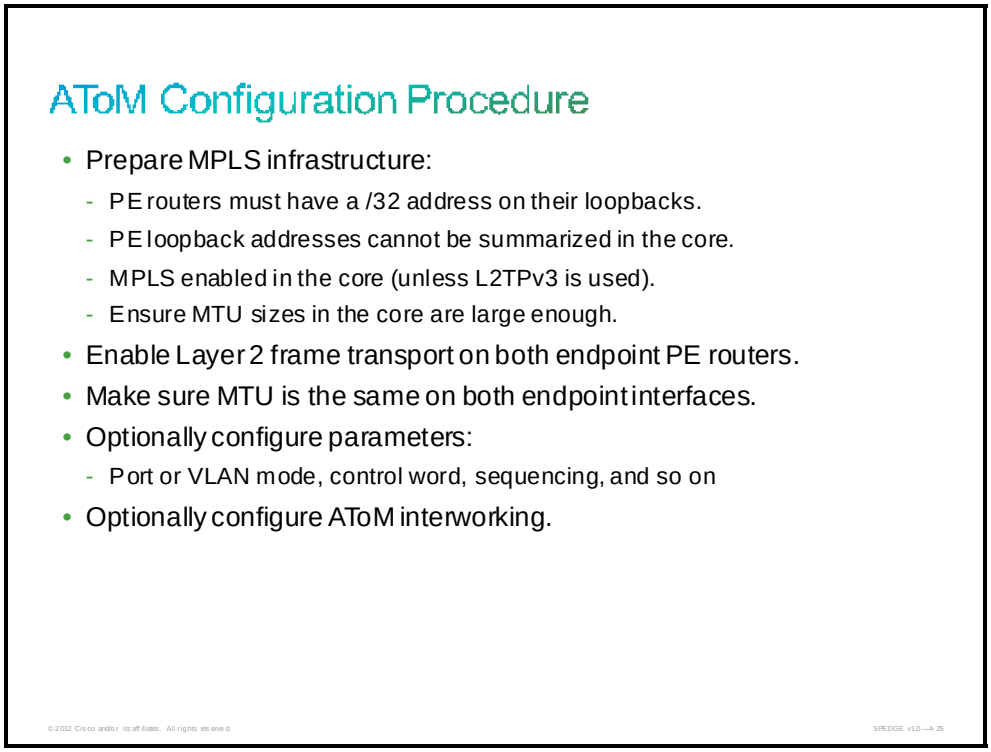


This diagram summarizes the EVC functionality by showing a physical Gigabit Ethernet interface with multiple EFP-based subinterfaces. The subinterfaces can be independently configured with flexible VLAN matching and VLAN tag rewriting. The incoming frames are then forwarded according to the defined parameters: linked to local attachment circuits, bridged or switched over a PW, or routed.

https://t.me/learningnets

AToM Implementation

This topic describes how to implement AToM.



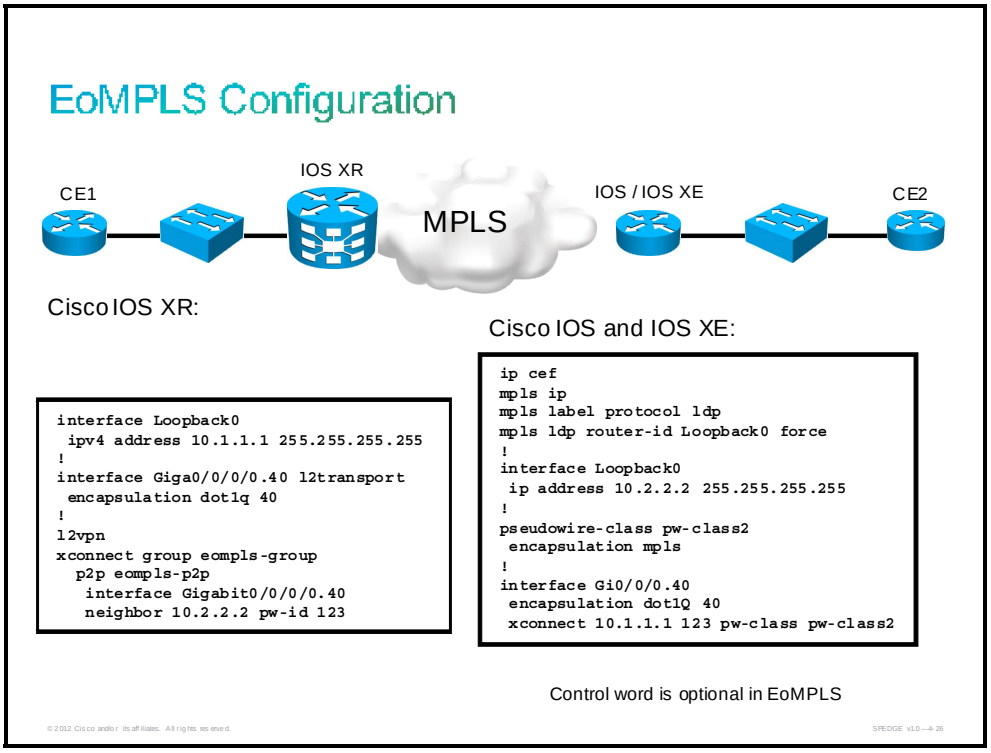
To implement AToM, you first need to prepare the MPLS infrastructure. This includes these steps:

- PE routers must have /32 addresses on their loopbacks. This is required for binding the VC label and the transport LDP into a label stack.
- The PE loopback addresses cannot be summarized in the core. Aggregation in the network would break the LSP path.
- MPLS must be enabled in the core unless L2TPv3 is used to provide IP-based transport.
- MTU sizes on the core links must be able to accommodate the customer MTU with the added label overhead.

Once the MPLS infrastructure is ready, you will do the following:

- Enable Layer 2 frame transport on both endpoint PE routers.
- Make sure MTU is the same on both endpoint interfaces.
- Optionally configure parameters, such as port or VLAN mode, the control word, sequencing, and so on.
- Optionally configure AToM interworking.

https://t.me/learningnets



EoMPLS is the most common AToM flavor. The configuration on Cisco IOS XR platforms differs from the one on Cisco IOS or IOS XE platforms.

On both systems, the pseudowire (PW) class is used as a container for optional parameters, such as encapsulation type, the use of control word, transport mode (port or VLAN), sequencing, and others.

On Cisco IOS XR Software, the attachment circuit must be enabled for Layer 2 transport using the **l2transport** keyword. The PW is configured using the **xconnect group** and **p2p** commands in circuit configuration mode. The **neighbor** command specifies the remote PE address and the VC ID.

On Cisco IOS and IOS XE platforms, the **xconnect** command in interface configuration mode defines the remote end, the VC ID, and, optionally, the PW class. On IOS XR platforms, the **xconnect** command is found in Layer 2 VPN configuration mode.

To specify a preferred interface for determining the LDP router ID, use the **mpls ldp router-id** command in global configuration mode.

The normal (default) method for determining the LDP router ID may result in a router ID that is not usable in certain situations. For example, an IP address that is selected as the LDP router ID might be an address that the routing protocol is not able to advertise to a neighboring router. The **mpls ldp router-id** command provides a means for specifying an interface whose IP address is to be used as the LDP router ID. The specified interface must be operational for its IP address to be used as the LDP router ID.

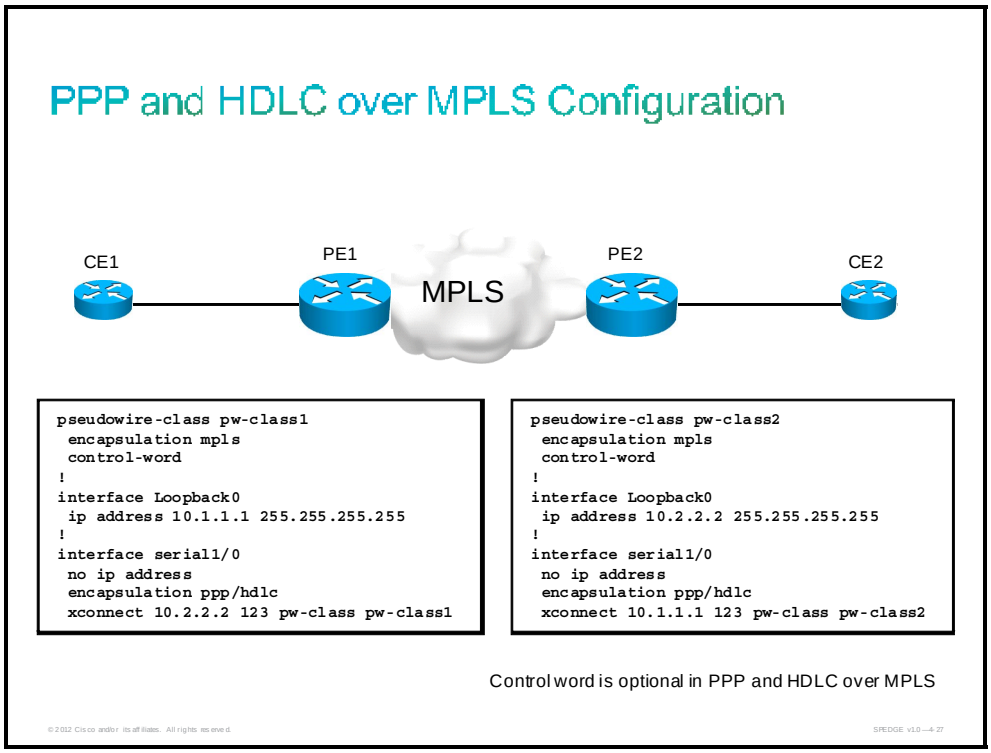
When it is executed without the **force** keyword, the **mpls ldp router-id** command modifies the method for determining the LDP router ID by causing the selection of the IP address of the specified *interface* argument (if the interface is operational) the next time that it is necessary to select an LDP router ID. The effect of the command is thus delayed until this need arises, which is typically the next time that the interface whose address is the current LDP router ID is shut down or the address itself is not configured.

https://t.me/learningnets

When it is executed with the **force** keyword, the effect of the **mpls ldp router-id** command depends on the current state of the specified interface:

- If the interface is up (operational) when the **mpls ldp router-id force** command is issued and if its IP address is not currently the LDP router ID, the LDP router ID is forcibly changed to the IP address of the interface. This forced change in the LDP router ID tears down any existing LDP sessions, releases label bindings that have been learned via the LDP sessions, and interrupts MPLS forwarding activity that is associated with the bindings.
- If the interface is down (not operational) when the **mpls ldp router-id force** command is issued, when the interface transitions to up, the LDP router ID is forcibly changed to the IP address of the interface. This forced change in the LDP router ID tears down any existing LDP sessions, releases label bindings that have been learned via the LDP sessions, and interrupts MPLS forwarding activity that is associated with the bindings

Note The VC ID, control word usage, and MTU sizes and PW type must match on the interfaces at both ends.



PPP and HDLC over MPLS is configured using the same commands as EoMPLS. This scenario shows Cisco IOS and IOS XE routers at both ends, because serial interfaces are more common for this router range.

The use of the control word is optional in PPP and HDLC over MPLS.

https://t.me/learningnets

Interworking Modes

- The attachment circuits are terminated locally.
- There are two types of interworking (any-to-any):
  - Ethernet (bridged):
    - Ethernet frames are extracted from attachment circuit and sent over PW.
    - VLAN tag is removed.
    - CEs can run Ethernet, BVI, or RBE.
    - Use the **interworking ip** command.
  - IP (routed):
    - IP packets are extracted from attachment circuit and sent over the PW.
    - Use the **interworking ethernet** command.

	AToM	L2TPv3	IP Mode	Ethernet
Frame Relay to Ethernet/VLAN	Yes	Yes	Yes	Yes
Frame Relay to PPP	Yes	Yes	Yes	No
Frame Relay to ATM AAL5	Yes	No	Yes	No
Ethernet/VLAN to ATM AAL5	Yes	No	Yes	Yes
Ethernet to VLAN	Yes	Yes	Yes	Yes

© 2012 Cisco and/or its affiliates. All rights reserved.

AToM interworking links two disparate Layer 2 encapsulations by terminating each attachment circuit locally and binding them over the PW.

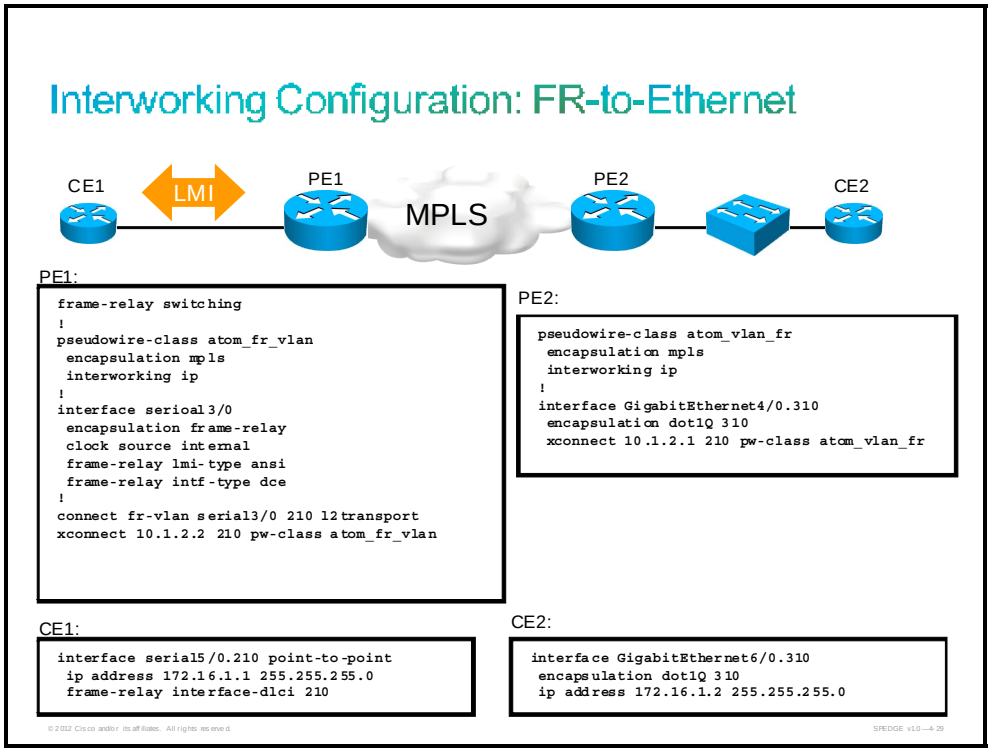
Interworking can be implemented using bridging or routing.

In bridging mode, the frames are extracted from the attachment circuit and sent over the PW. In the case of Ethernet, the VLAN tag (if available) is removed. In that case, the CEs can run Ethernet, bridge-group virtual interface (BVI), or routed bridge encapsulation (RBE).

In routing mode, the IP packets are extracted from the attachment circuit and sent over the PW to the remote end.

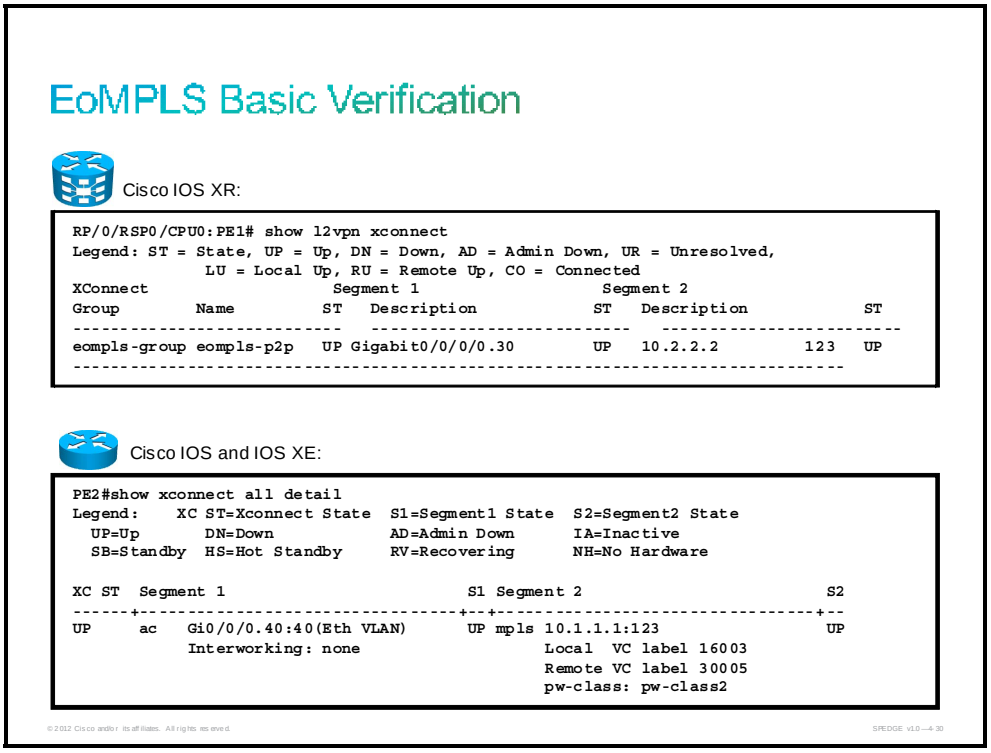
The figure lists various interworking methods and their feasibility in AToM, L2TPv3, routing, and bridging mode.

https://t.me/learningnets



These configurations depict a Frame Relay-to-Ethernet interworking scenario using routing mode. The mode is defined with the **interworking ip** command in the pseudowire-class configuration mode.

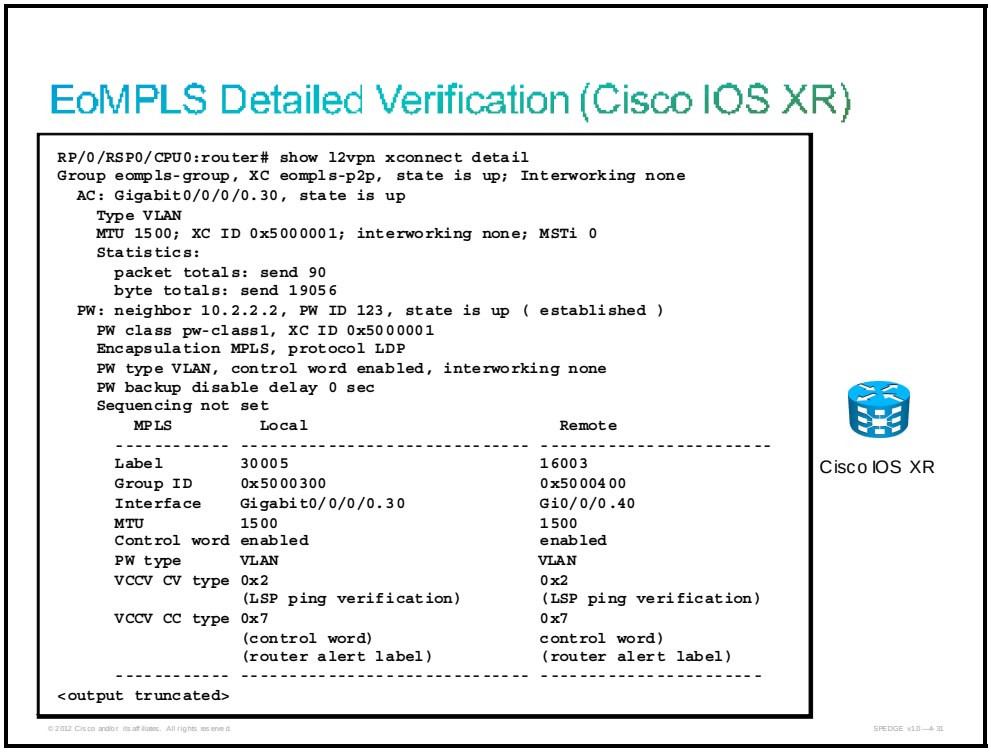
https://t.me/learningnets



The EoMPLS operation is verified using different commands on Cisco IOS XR and Cisco IOS and IOS XE platforms. The **show l2vpn xconnect** command provides brief information on configured cross-connects.

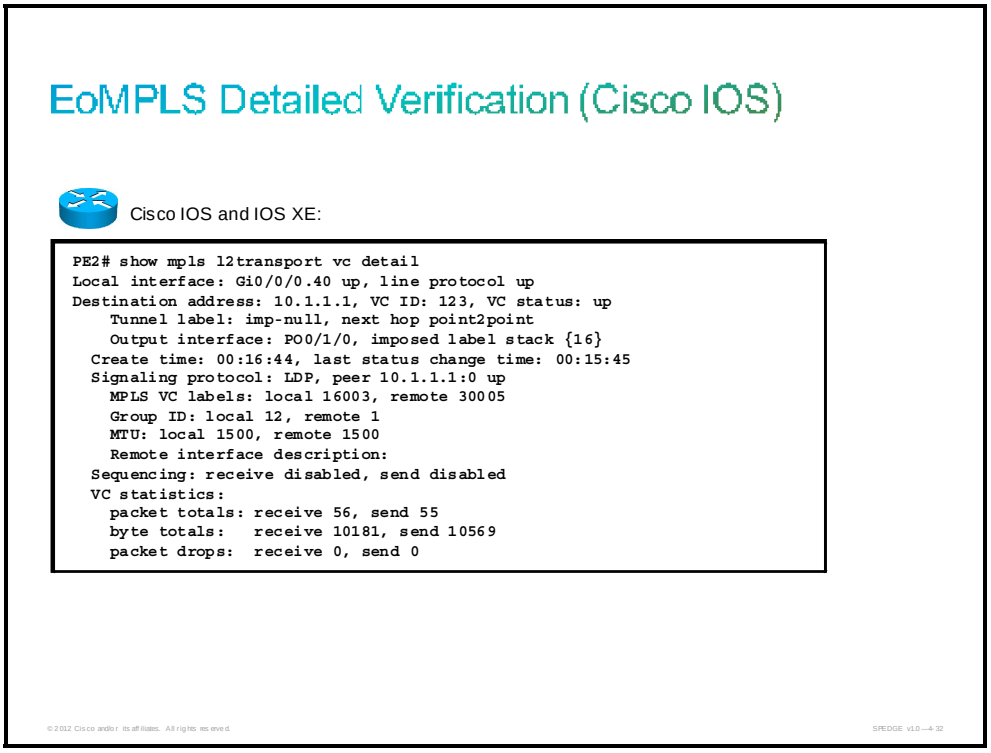
On Cisco IOS routers, you can display equivalent information with the **show xconnect all detail** command.

https://t.me/learningnets



Cisco IOS XR devices offer a much more granular output when you use the **show l2vpn xconnect** command. The output includes information about the local (configured) and remote (signaled) parameters, such as labels, group ID, attachment circuit interface, MTU size, control word usage, PW type (port or VLAN), and VCCV connectivity verification and control channel types. If any of these parameters do not match on both sides, the PW does not reach the up state.

https://t.me/learningnets

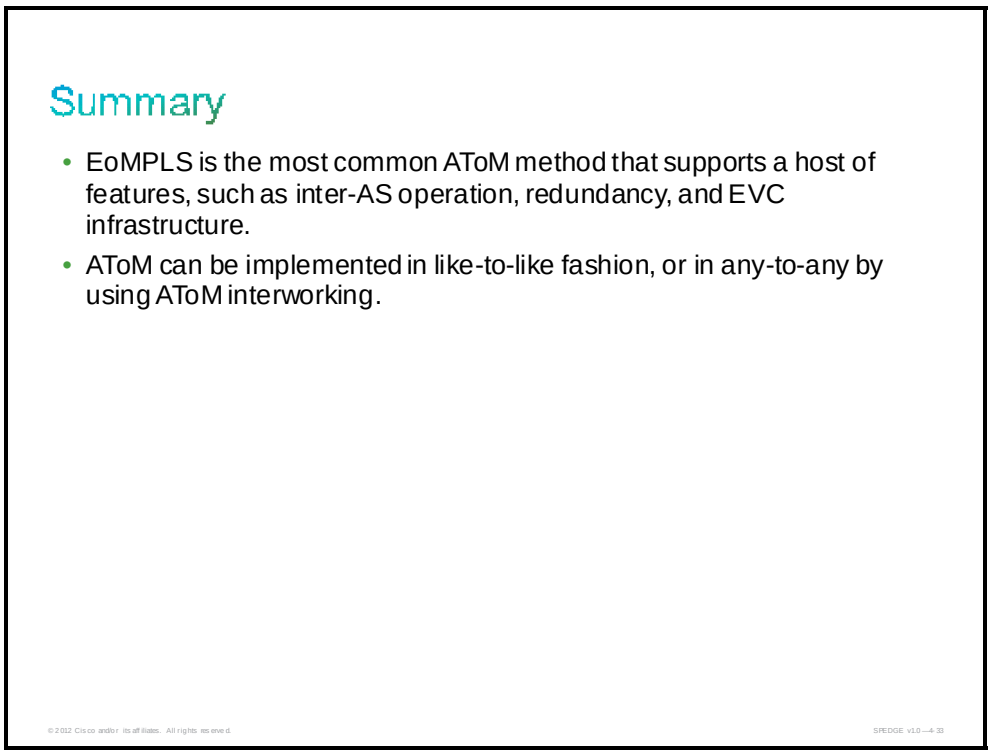


In Cisco IOS Software, you can investigate the VC-related information by issuing the **show mpls l2transport vc detail** command. It displays the local and remote interfaces, exchanged labels, group ID, MTU sizes, sequencing status, and VC statistics.

https://t.me/learningnets

Summary

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



https://t.me/learningnets

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 3

Implementing VPLS

Overview

Virtual Private LAN Services (VPLS) can be used to implement Carrier Ethernet in a service provider MPLS infrastructure.

This lesson describes the VPLS implementation on Cisco IOS XR routers.

Objectives

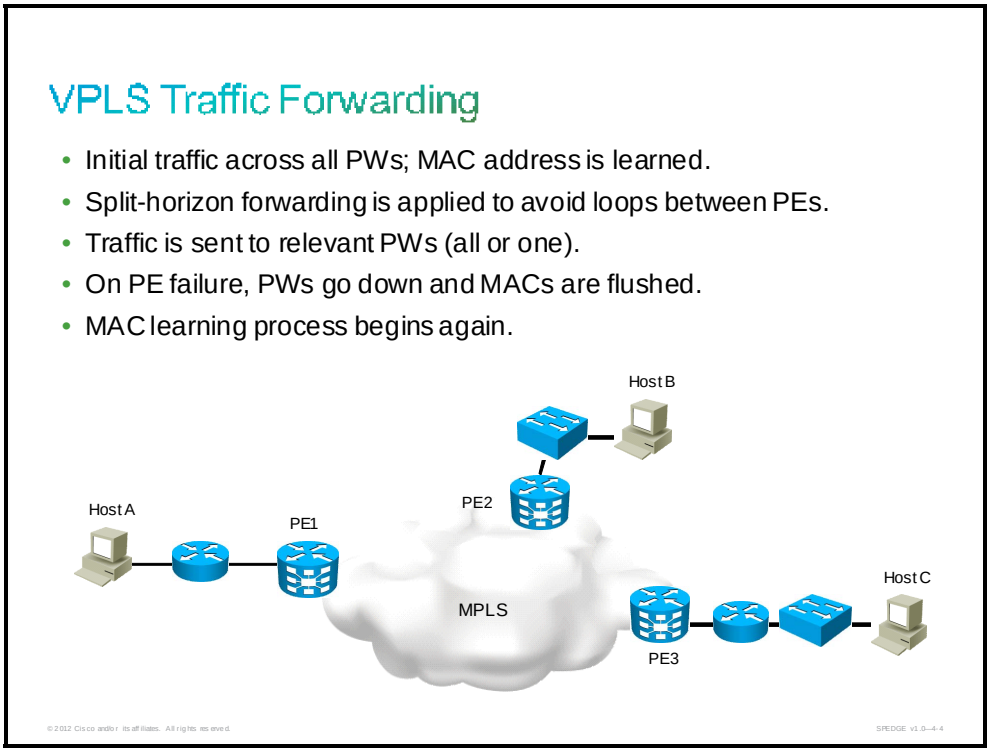
Upon completing this lesson, you will be able to describe Ethernet services that are used in the service provider network. You will be able to meet these objectives:

- Discuss VPLS
- Implement VPLS and H-VPLS

https://t.me/learningnets

VPLS Overview

This topic discusses the VPLS implementation technique.



Forwarding of Unicast Frames

Suppose that the MAC address of host C is C and the MAC address of host B is B, for the customer network in the figure. Suppose host C sends a frame with source MAC address C and destination MAC address B. Suppose that PE3 does not know the location of MAC address B. As a learning bridge would do, PE3 floods the packet on all ports except the port on which it arrived. This means the packet is flooded to the pseudowire (PW) to PE2 and the PW to PE1.

PE1 and PE2 know that the packet belongs to the customer VPLS, by virtue of the PW on which the frame arrived. PE1 and PE2 both perform destination MAC address lookup in their VPLS forwarding tables corresponding to this customer. If PE1 does not know the location of MAC address B, it floods the frame on its local ports to the customer edge (CE). However, it does not flood the frame to any other provider edge (PE) routers. This split-horizon scheme ensures no forwarding loops occur. Similarly, PE2 forwards the frame to the port facing the switch CE.

Receiving frames with MAC address C enables each PE to learn the location of host C. Thus, PE2 and PE1 create an entry in their forwarding tables with an association between MAC address C and their respective PWs to PE3. In this way, all PEs learn the MAC addresses and create an association between MAC addresses and PWs (for remote destinations) in their forwarding tables for that particular VPLS instance.

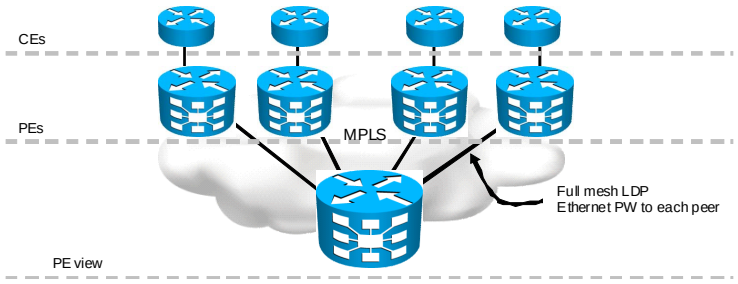
Forwarding of Broadcast and Multicast Frames

Suppose PE3 receives a broadcast frame sent by host C. The frame must be sent to all sites of the customer VPLS. PE3 floods the frame on PWs to PE1 and PE2. In turn, PE1 and PE2 flood the frame to the attached CEs, but due to split horizon, do not send the frames to any PEs. Multicast traffic is also treated this way.

https://t.me/learningnets

VPLS Topology: PE View

- Each PE has a point-to-multipoint view of all other PEs:
  - Sees itself as a root bridge with split-horizon loop protection
- Full mesh topology obviates STP in the service provider network.
- Customer STP is transparent to the service provider:
  - Customer BPDUs are forwarded transparently.

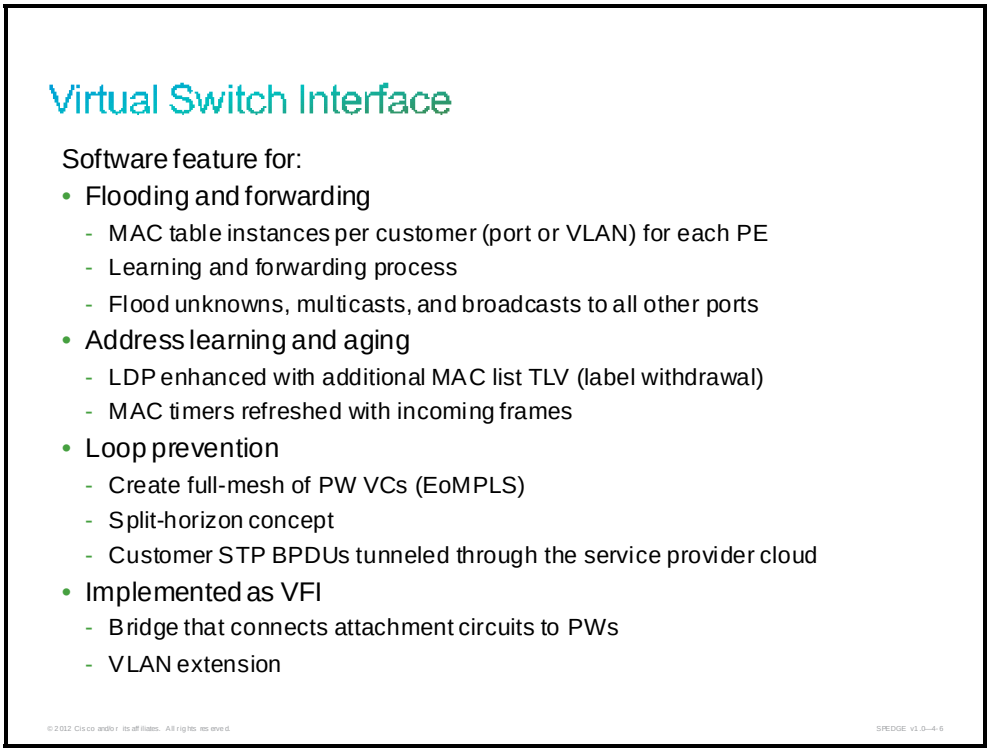


© 2012 Cisco and/or its affiliates. All rights reserved.

The full mesh of PWs in the service provider network allows the implementation of the split-horizon principle that is similar to the split horizon of an internal BGP mesh. Frames received over a PW are generally not forwarded to another PW. This concept provides loop-free forwarding. Therefore, there is no need for Spanning Tree Protocol (STP) in the service provider network. The customer STP bridge protocol data units (BPDUs) are transparently tunneled over the PWs to detect loops on the customer layer.

Note	The split-horizon rule must be disabled in certain environments, such as hierarchical VPLS (H-VPLS). H-VPLS is described later in the lesson.

https://t.me/learningnets



The Virtual Switch Interface (VSI), implemented as a virtual forwarding instance (VFI), is a virtual Layer 2 Forwarding entity that defines the VPLS domain membership and resembles virtual switches on PE routers. A VPLS domain consists of Ethernet interfaces or VLANs that belong to the same (virtual) LAN but are connected to multiple PE devices. The VSI learns remote MAC addresses and is responsible for proper forwarding of the customer traffic to the appropriate end nodes. It is also responsible for guaranteeing that each VPLS domain is loop-free. The VSI is responsible for several functions, namely MAC address management, dynamic learning of MAC addresses on physical ports and virtual circuits (VCs), aging of MAC addresses, MAC address withdrawal, flooding, and data forwarding.

All PWs in a VFI are placed by default into the same split-horizon group, which effectively prevents traffic from forwarding to other PWs in the same VFI.

With VPLS Layer 2 VPNs, the customer can connect with a switch or a router.

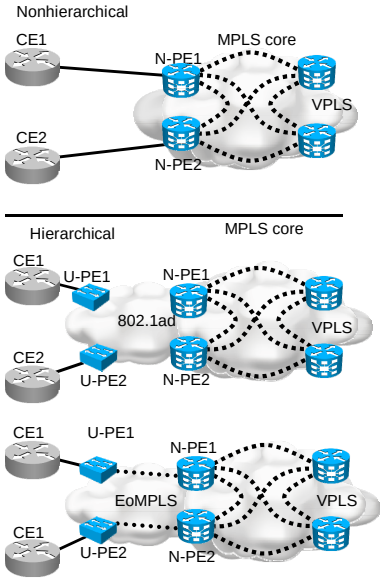
If connecting with a router, the VPLS just looks like a switch to the routing protocols on each side. No MAC learning is done beyond the MAC address of the directly connected CE router interface.

If connecting at Layer 2 with a switch, then MAC learning on the PE of all the MAC addresses on the directly connected customer LAN is possible. This is a situation where MAC limiting or another filtering method would be implemented, or a provider backbone bridge (PBB) (802.1ah) can be used.

https://t.me/learningnets

VPLS Architectures

- Two architectures:
  - Nonhierarchical
    - Single PE (flat)
  - Hierarchical (H-VPLS)
    - With Ethernet access
      - 802.1ad (IEEE standard for QinQ)
    - With MPLS access
- Two PE roles:
  - Network-facing PE (N-PE)
    - VPLS termination
    - Layer 3 services
  - User-facing PE (U-PE)
    - Customer UNI



© 2012 Cisco and/or its affiliates. All rights reserved.

Hierarchical VPLS (H-VPLS) topologies have been developed to address the scaling limitations within flat VPLS architecture. VPLS requires a full mesh of tunnel LSPs between all PE routers that participate in the VPLS service. For each VPLS service, n*(n-1)/2 PWs must be set up between PEs. This creates signaling overhead. The packet replication requirement deters large-scale deployment. Hierarchical connectivity reduces signaling and replication overhead.

There are two access mechanisms in H-VPLS: Ethernet and MPLS. Ethernet-based access uses the IEEE 802.1ad standard or its predecessor, 802.1QinQ.

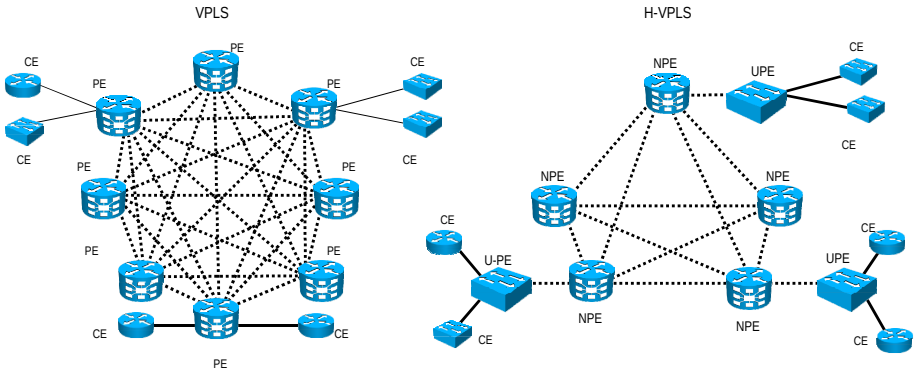
H-VPLS defines user-facing PE (UPE) routers and network-facing PE (NPE) routers.

https://t.me/learningnets

VPLS vs. H-VPLS

- Significant signaling overhead
- Full PW mesh from the edge
- Node discovery and provisioning extends end to end

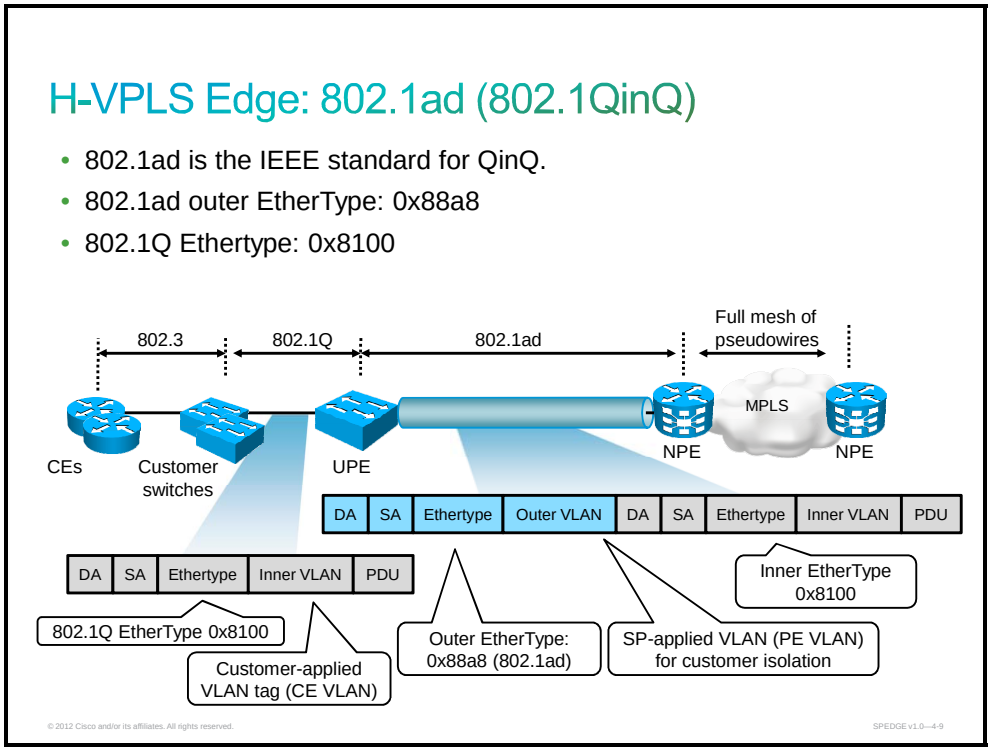
- Minimizes signaling overhead
- Full PW mesh among core devices
- Partitions node discovery process



© 2012 Cisco and/or its affiliates. All rights reserved.

The H-VPLS hierarchy provides the benefits of less signaling in an MPLS core network and less packet replication on NPE routers. The UPE routers have an aggregation role and do some packet replication and MAC address learning.

https://t.me/learningnets



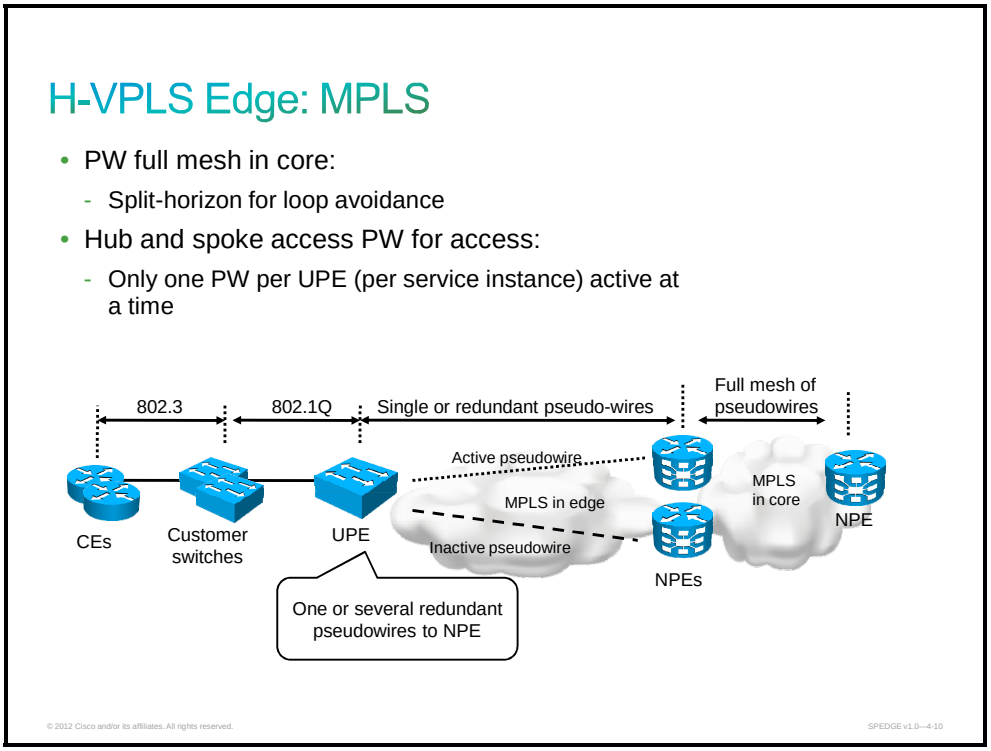
IEEE 802.1QinQ is an Ethernet networking standard formally known as IEEE 802.1ad. It is also known as provider bridging, stacked VLANs, or queue-in-queue (QinQ).

802.1QinQ specifies architecture and bridge protocols to provide separate instances of the MAC services to multiple independent users in a manner that does not require cooperation among the users, and requires a minimum of cooperation between the users and the provider. The idea is to provide the possibility for customers to run their own VLANs inside the VLAN of a service provider. This way, the service provider can just configure one VLAN for the customer, and the customer can then treat that VLAN like a trunk.

IEEE 802.1QinQ has these characteristics:

- 802.1Q has a 12-bit VLAN ID field, which has a theoretical limit of 2^12=4096 tags. With the growth of networks, this limitation has become more acute. A double-tagged frame has a theoretical limitation of 4096*4096=16,777,216, sufficient to accommodate network growth.
- The addition of a second tag allows new set of operations. Having multiple tags—the tag stack—allows switches to more easily modify frames. In a tag stack scheme, switches can add (“push”), remove (“pop”) or modify single or multiple tags.
- A multitagged frame not only has multiple VLAN IDs, but has multiple EtherTypes and other VLAN header bit fields.
- A tag stack creates a mechanism for Internet service providers to encapsulate customer single-tagged 802.1Q traffic with a single tag, the final frame being a QinQ frame. The outer tag is used to identify and segregate traffic from different customers; the inner tag is preserved from the original frame.
- 802.1QinQ is upward compatible with 802.1Q. Although 802.1QinQ is limited to two tags, there is no ceiling on the standard limiting a single frame to more than two tags, allowing for growth in the protocol. In practice, service provider topologies often anticipate and utilize frames with more than two tags.

https://t.me/learningnets



The second approach uses MPLS in the edge network, instead of 802.1ad, which was described in the previous method.

In the MPLS-based edge, the VPLS core PWs (hub) are augmented with access PWs (spoke) to form a two-tier H-VPLS. The customer sites are connected to UPE devices. The UPE devices have single or redundant connection (PW) to NPE routers. The NPE routers are connected in a basic VPLS full mesh service. For each VPLS service, a single spoke PW is set up between UPE and NPE devices. Unlike traditional PWs that terminate on a physical or logical port, a spoke PW terminates on a VSI on UPE and NPE devices. The UPEs and NPEs treat each spoke connection like an attachment circuit of the VPLS service. The PW label is used to associate the traffic from the spoke PW to a VPLS instance.

UPE Operation

The UPE device supports Layer 2 switching and does all the normal bridging functions of learning and replication on all its ports, including the spoke PW. Packets to an unknown destination are replicated to all ports in the service, including the spoke PW. Once the MAC addresses of CE devices connected to the same UPE device are learned, traffic between them is switched locally, saving the capacity of the spoke PW to NPEs. Similarly, traffic between remotely connected CE devices to different UPE devices is switched directly onto the spoke PW and sent to NPEs over the point-to-point PW. Since the UPE is bridging-capable, only a single PW is required per VPLS instance for any number of access connections in the same VPLS service. This reduces the signaling overhead between UPEs and NPEs.

NPE Operation

An NPE device supports all bridging functions for VPLS service and supports the routing and MPLS encapsulation as in basic VPLS. The operation of the NPE is independent of the type of device at the other end of the spoke PW. Thus the NPE will switch traffic between the spoke PW, hub PWs, and attachment circuits once it has learned the MAC addresses.

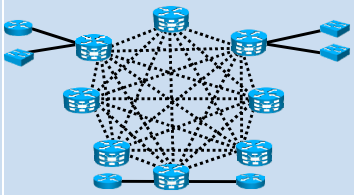
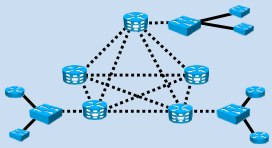
https://t.me/learningnets

Dual-Homed UPE

The failure of an NPE router or the connection of the PW can cause the UPE to suffer a total loss of connectivity. To prevent this, redundant connections can be provided. The UPE is dual-homed into two NPE routers. One of the two PWs is designated as primary and the other as standby. The UPEs negotiate PW labels for both PWs, but do not use the standby PW unless the primary PW fails. Spanning-tree instance or manual configuration can be used to designate primary and standby status.

Upon failure of the primary PW, the UPE immediately switches to the standby PW. At this point, the NPE that terminates the standby PW starts learning MAC addresses on the spoke PW. All other NPEs initially continue to send traffic to the initial NPE until they learn that the devices are now connected to a new NPE. To enable a faster unlearning process, the new NPE may send out a flush message using the MAC list type, length, and value (TLV) (type 0x404) to all NPEs. Upon receiving the message, the NPEs flush the MAC addresses associated with that VPLS instance.

VPLS and H-VPLS Evaluation

	Flat VPLS	H-VPLS
		
Pros	• Simple provisioning	• Suitable for large environments • Reduced replication and signaling overhead on NPEs • Expansion affects new nodes only
Cons	• Scalability limitation to small environments • PE packet replication • Directed LDP full mesh – n * (n-1)/2 sessions	• More complicated provisioning • More complex design and operations • More expensive hardware for MPLS-based access

© 2012 Cisco and/or its affiliates. All rights reserved.

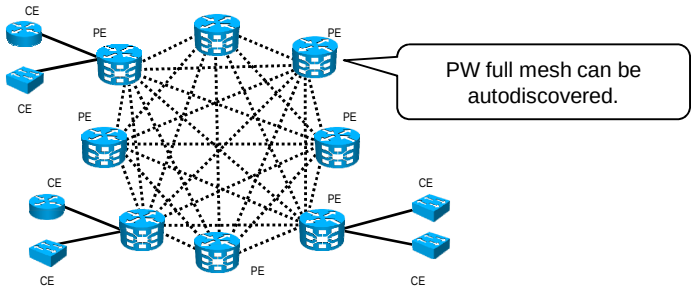
Flat VPLS offers the easiest implementation method. It supports both port and VLAN transport modes, but is limited to small deployments. All packet replication is performed on the PE devices. It results in a directed LDP full mesh of n*(n-1)/2 sessions.

H-VPLS is suitable for larger environments. It places less replication burden and signaling overhead on the NPE routers. It allows the service provider to expand certain network areas without affecting the core and most of the NPE infrastructure. It requires, however, more complex design, provisioning, and operations, and may incur higher cost, if MPLS is deployed in the edge.

https://t.me/learningnets

VPLS Neighbor Autodiscovery

- Requires special address family for MP-BGP neighbors
  - **address-family l2vpn vpls-vpws**
- Available for VPLS and VPWS
- Two signaling methods: LDP and BGP
  - Both methods use VLAN IDs, RDs, and RTs to limit discovery scope.



© 2012 Cisco and/or its affiliates. All rights reserved.

Once the PEs have ascertained that other PEs have an association with the same VPLS instance, each PE needs to set up a PW between themselves and bind the PWs to the particular VSI. These PWs can be set up using one of these methods:

For neighbor autodiscovery:

- Static configuration
- BGP-based autodiscovery

For PW signaling:

- BGP-based signaling
- LDP-based signaling

LDP-Based Signaling

LDP is used for the signaling of PWs that are used to interconnect the VPLS instance of a given customer on the PEs. In order to signal a full mesh of PWs, a full mesh of Targeted LDP (T-LDP) sessions is required between the PEs. In the absence of autodiscovery, these sessions must be manually configured on each PE router. This LDP session is used to communicate the inner label or VC label that must be used for each PW. The network operator assigns a VC ID that is used to identify a particular PW, and it is configured to be the same for a particular VPLS instance on all PE routers.

The LDP session communicates the label value and the FEC element that includes the VC ID, control word bit (indication whether a control word will be used), VC type (Ethernet or VLAN-tagged Ethernet), and interface parameters field (media maximum transmission unit (MTU), and so on.)

https://t.me/learningnets

LDP signals the setup, maintenance, and teardown of a PW between PE devices. Once a PE discovers that other PEs have an association for a particular VPLS instance, the PEs signal using T-LDP to other PEs that a PW is required to be set up between PEs. When a PE device is associated with a particular VSI, LDP transmits a label-mapping message with VC type 0x0005 and a 4-byte VC ID value. If the remote PE has an association with that particular VC ID, it will accept the LDP label-mapping message and respond with its own label-mapping message. Once the two unidirectional VCs are operational, they are combined to form a bidirectional PW.

BGP-Based Autodiscovery and Signaling

The BGP Network Layer Reachability Information (NLRI) takes care of autodiscovery and signaling at the same time. The NLRI is generated by a given PE containing the necessary information required by any other PE. These components enable the automatic setting up of a full mesh of PWs for each VPLS.

On each PE, a route distinguisher (RD) and a route target (RT) are configured for each VPLS, like in Layer 3 VPN and Layer 2 VPN. The RT is same for a particular VPLS across all PEs, and is used to identify which VPLS a particular BGP message pertains to. The RD is used to disambiguate routes. On each PE, for each VPLS, an identifier is configured called a VPLS Edge Identifier (VE ID). Each PE involved in a particular VPLS must be configured with a different VE ID. BGP is used to advertise the VE ID to other PEs. This, along with other information in NLRI, is used to calculate the value of the PW label required to reach the advertising PE.

For a given VPLS, a PE requires that each remote PE uses a different PW label to send traffic to that PE. This requirement is done to facilitate the MAC learning process. This way, the receiving PE can learn which PE is associated with the source MAC address of the frame. A PE could send a list of PW labels that are required to reach it, one per remote VPLS edge in that VPLS. However, this would mean that a PE sends a long list of labels if there is a large number of PEs. Instead, the necessary information is carried in Border Gateway Protocol (BGP) NLRI. This allows all remote PEs to calculate the PW label expected by the advertising PE.

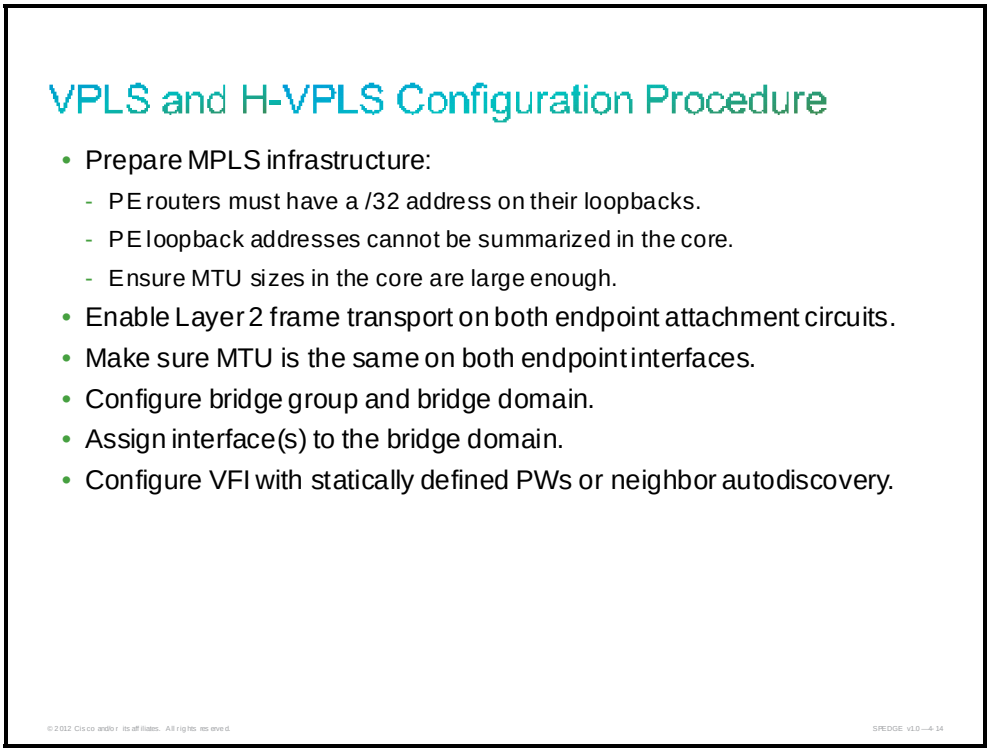
A BGP update message contains the extended community RT (which allows the receiving PE to ascertain which particular VPN the advertisement pertains to), Layer 2 information (control flags to indicate whether a control word is included, the encapsulation type—Ethernet or VLAN-tagged—and MTU), and other BGP attributes such as autonomous system (AS) path, origin, and so on.

The NLRI contains the RD, VE ID, label base, VE block offset, and VE block size. The label base, VE block offset and VE block size are the information that the receiving PE requires to calculate the PW label when sending traffic for VPLS to the advertising PE.

https://t.me/learningnets

Implementing VPLS and H-VPLS

This topic describes the implementation of VPLS and H-VPLS.



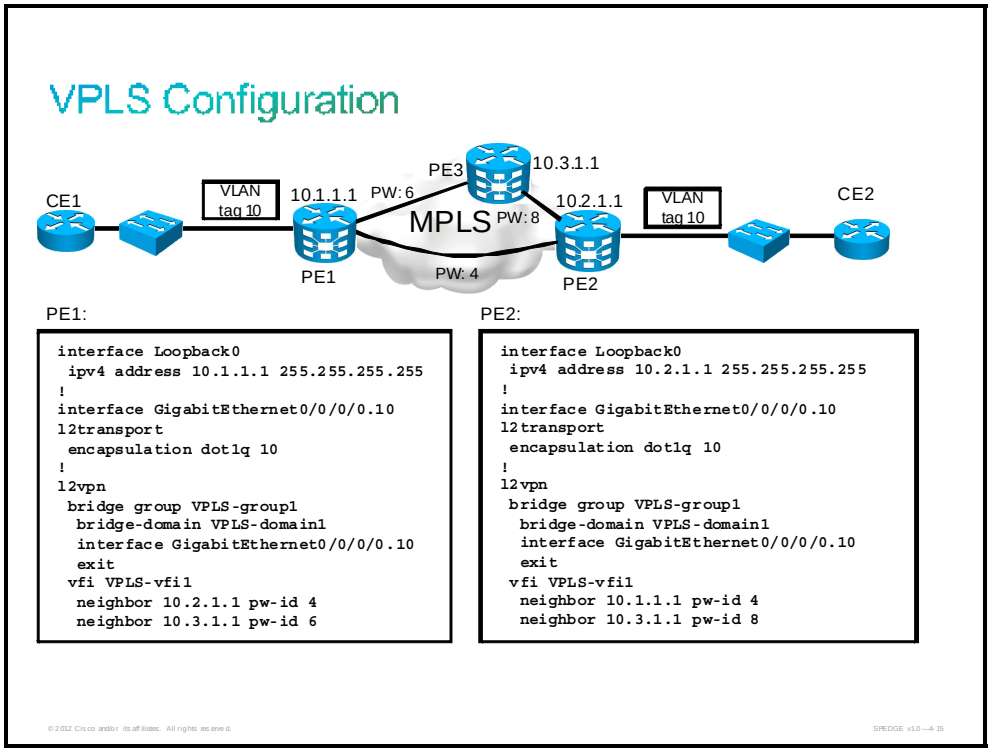
To implement VPLS or H-VPLS, you first need to prepare the MPLS infrastructure. This process includes these steps:

- PE routers must have a /32 address on their loopbacks. This is required for binding the VC label and the transport LDP into a label stack.
- PE loopback addresses cannot be summarized in the core. Aggregation in the network would break the LSP path.
- MPLS must be enabled in the core.
- MTU sizes on the core links must be able to accommodate the customer MTU with the added label overhead.

Once the MPLS infrastructure is ready, you will do the following:

- Enable Layer 2 frame transport on both endpoint PE routers.
- Make sure the MTU is the same on both endpoint interfaces.
- Configure the bridge group and bridge domain.
- Assign interfaces to bridge domain.
- Configure the VFI with either statically defined PWs or neighbor autodiscovery.

https://t.me/learningnets



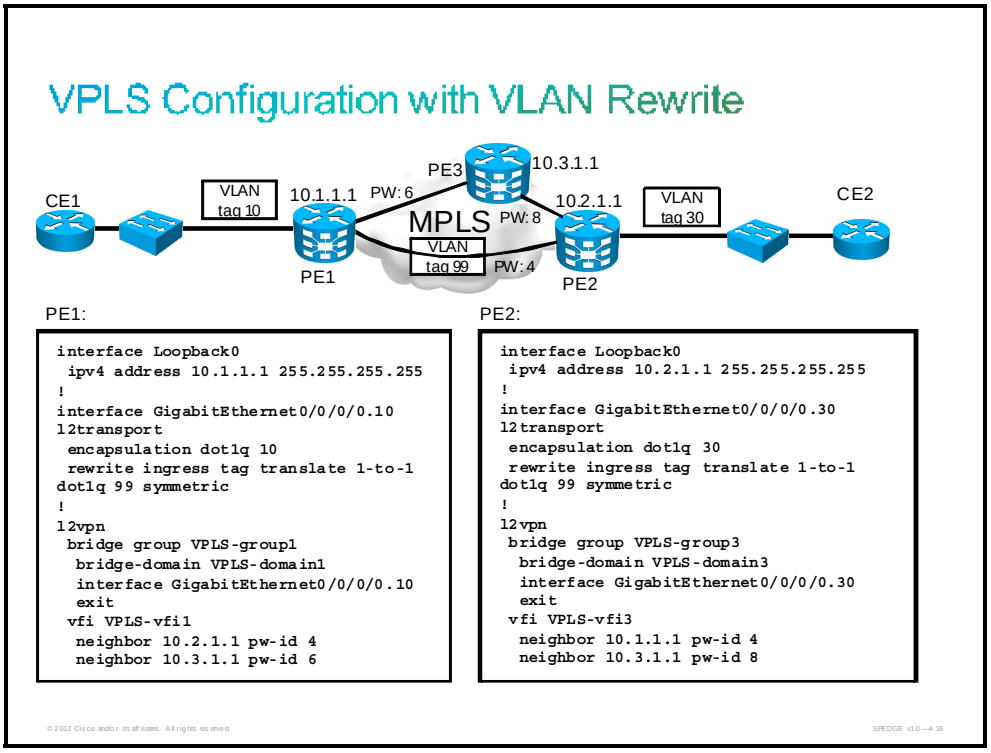
The first scenario is a flat VPLS implemented on Cisco IOS XR platforms.

The attachment circuits on both ends are enabled for Layer 2 transport using the **l2transport** keyword.

The VPLS is configured within the Layer 2 VPN (**l2vpn** command), bridge group, and bridge domain configuration mode. The bridge domain contains two main configuration elements: a list of the local attachment circuits, and the VFI that contains all required PWs.

The names of the bridge group and domain have local significance. The PW ID must match on both ends.

https://t.me/learningnets

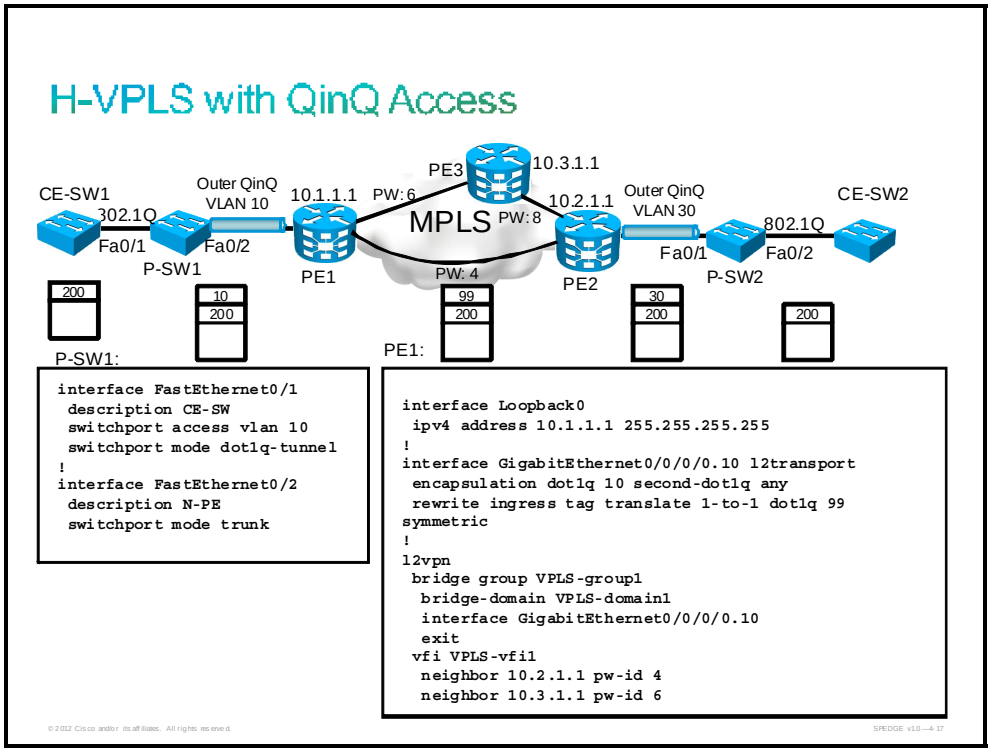


This scenario differs from the previous one in managing the outmost VLAN tag. In this example, the VLAN tag is replaced with the VLAN tag 99. On the egress PE, the VLAN tag 99 will be swapped with the local VLAN tag 30.

The service provider will use rewrite because the service provider offers some service on VLAN 99, and customers uses different VLANs.

This technique illustrates how the service provider can replace customer VLAN tags with a tag from its own range for transport across the core network.

https://t.me/learningnets



H-VPLS requires that either QinQ or MPLS is used in the access network. This scenario depicts QinQ access.

The switch must be configured for QinQ tunneling. This requires the **switchport mode dot1q-tunnel** and the **switchport access vlan** commands on the CE-facing port. The switchport **mode dot1q-tunnel** command instructs the switch to encapsulate incoming 802.1Q frames within the VLAN tag that is specified with the **switchport access vlan** command. In this case, P-SW1 applies the outer VLAN tag 10 to all frames that are received from the customer switch CE-SW1.

The PE routers are configured for regular VPLS with static or autodiscovered PWs. The encapsulation on PE1 is set using the command **encapsulation dot1q 10 second-dot1q any** to match only QinQ frames and exclude 802.1Q frames. The command **encapsulation dot1q 10** would also work in this scenario. In this case, PE1 translates the outer VLAN tag to 99 for transport to the egress PE.

PE2 receives the VPLS packets and translates the outer VLAN tag 99 to the VLAN tag allocated to the local customer, 30. This is achieved with the **rewrite ingress tag translate 1-to-1 dot1q 99 symmetric** command along with the **encapsulation dot1q 30 second-dot1q any** command.

The configurations of PE2 and P-SW2 are configured in the same way, symmetrically to PE1 and P-SW1, except that the customer frames are encapsulated in VLAN tag 30 instead of 10.

https://t.me/learningnets

Verifying VPLS Operation

```
RP/0/RSP0/CPU0:PE3# show l2vpn bridge-domain detail
Sat Nov 26 13:48:47.127 UTC
Bridge group: VPLS-group3, bridge-domain: VPLS-domain3, id: 1, state: up, ShgId: 0,
MSTi: 0
  MAC learning: enabled
  MAC withdraw: enabled
    MAC withdraw for Access PW: enabled
  Flooding:
    Broadcast & Multicast: enabled
    Unknown unicast: enabled
  MAC aging time: 300 s, Type: inactivity
  MAC limit: 4000, Action: none, Notification: syslog
  MAC limit reached: no
  MAC port down flush: enabled
  MAC Secure: disabled, Logging: disabled
  Split Horizon Group: none
  Dynamic ARP Inspection: disabled, Logging: disabled
  IP Source Guard: disabled, Logging: disabled
  DHCPv4 snooping: disabled
  IGMP Snooping profile: none
  Bridge MTU: 1500
  MIB cvplsConfigIndex: 2
  Filter MAC addresses:
  Create time: 26/11/2011 11:38:38 (02:10:08 ago)
  No status change since creation
  ACs: 1 (1 up), VFIs: 1, PWs: 1 (1 up), PBBs: 0 (0 up)
  < to be continued>
```

© 2012 Cisco and/or its affiliates. All rights reserved.

The **show l2vpn bridge-domain detail** command is used to verify the VPLS operation. It provides a verbose output that is only partly shown here. This section describes parameters that relate to bridging, MAC learning, and security features active with the particular VPLS domain

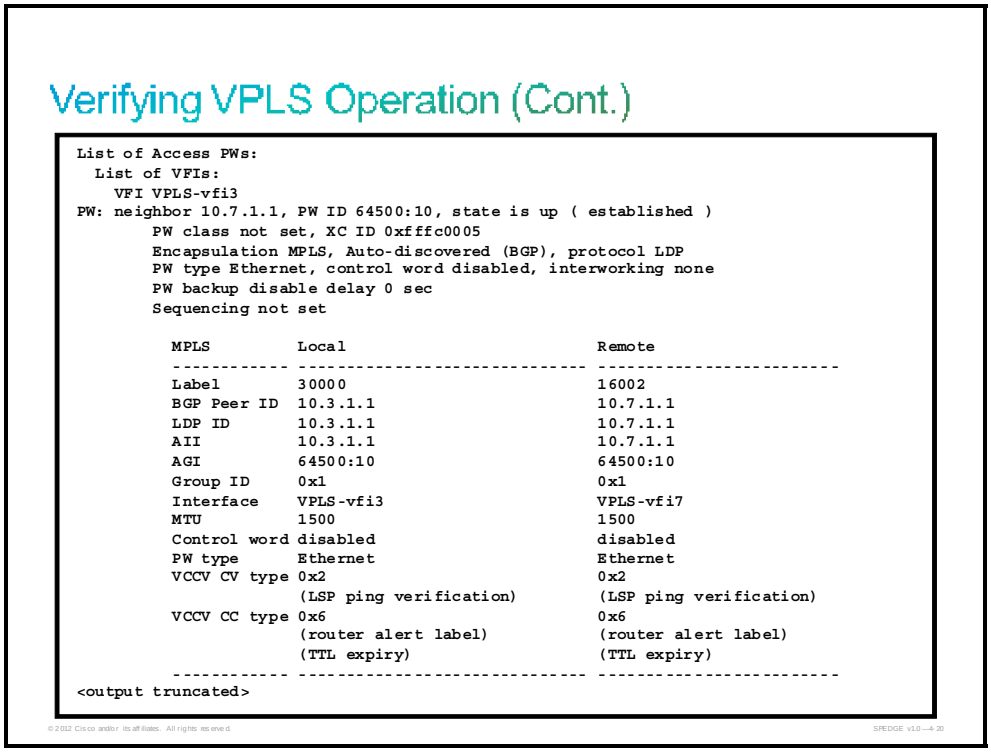
Verifying VPLS Operation (Cont.)

```
List of ACs:
    AC: GigabitEthernet0/0/0/0.30, state is up
      Type VLAN; Num Ranges: 1
      VLAN ranges: [30, 30]
      MTU 1504; XC ID 0x840001; interworking none
      MAC learning: enabled
      Flooding:
        Broadcast & Multicast: enabled
        Unknown unicast: enabled
      MAC aging time: 300 s, Type: inactivity
      MAC limit: 4000, Action: none, Notification: syslog
      MAC limit reached: no
      MAC port down flush: enabled
      MAC Secure: disabled, Logging: disabled
      Split Horizon Group: none
      Dynamic ARP Inspection: disabled, Logging: disabled
      IP Source Guard: disabled, Logging: disabled
      DHCPv4 snooping: disabled
      IGMP Snooping profile: none
      Storm Control: disabled
      Static MAC addresses:
      Statistics:
        packets: received 31686, sent 27420
        bytes: received 2156476, sent 1911176
      Storm control drop counters:
        packets: broadcast 0, multicast 0, unknown unicast 0
        bytes: broadcast 0, multicast 0, unknown unicast 0
<to be continued>
```

© 2012 Cisco and/or its affiliates. All rights reserved.

The next part of the output describes all attachment circuits assigned to the bridge domain. It provides information on enabled security features and various statistics.

https://t.me/learningnets

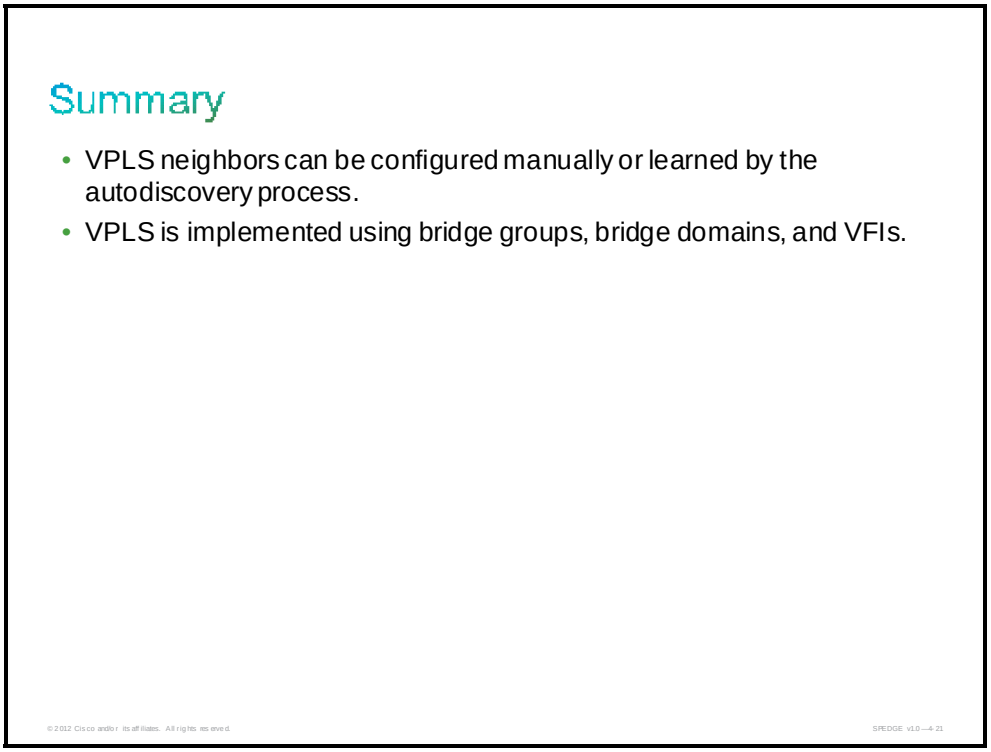


The last major part of the output focuses on the PWs that constitute the VFI. In this section, you can view the PW parameters on the local side and on the remote end.

https://t.me/learningnets

Summary

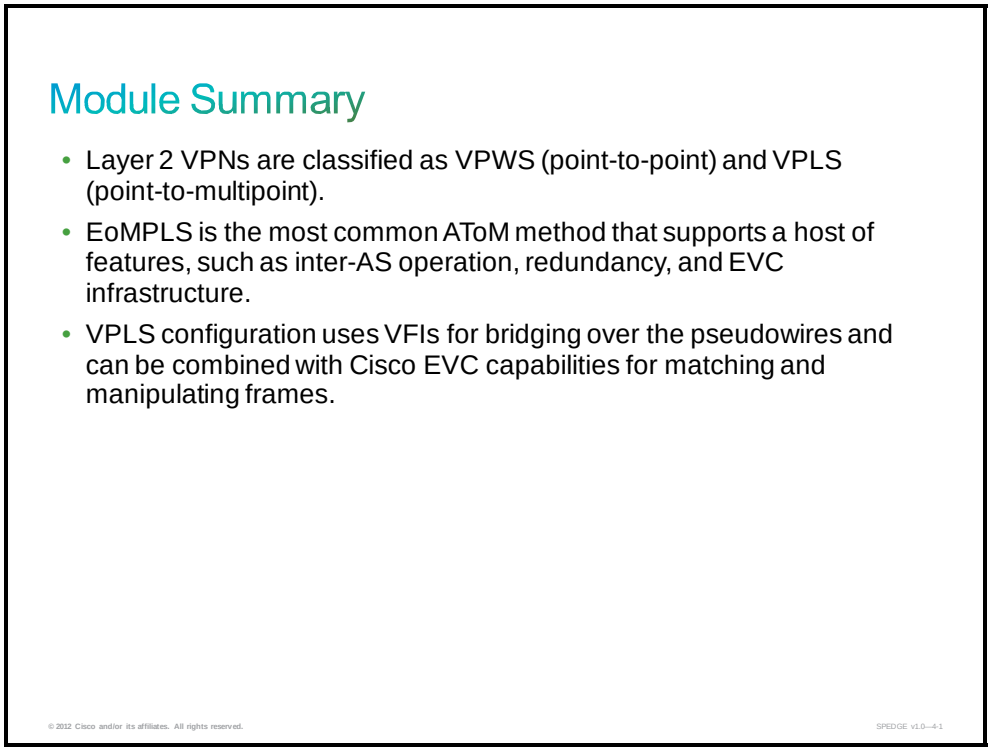
This topic summarizes the primary points that were discussed in this lesson.



https://t.me/learningnets

Module Summary

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



Layer 2 VPNs are classified as Virtual Private Wire Service (VPWS) (point-to-point) and Virtual Private LAN Services (VPLS) (point-to-multipoint). They can be transported over Multiprotocol Label Switching (MPLS) or IP networks. VPLS can only be carried over MPLS and cannot use Layer 2 Tunneling Protocol version 3 (L2TPv3), which provides the transport protocol in the pure IP core.

Any Transport over MPLS (AToM) is a subset of VPWS. It is deployed to carry Ethernet, Frame Relay, PPP and High-Level Data Link Control (HDLC) frames, and ATM cells, and it emulates time-division multiplexing (TDM) circuits over MPLS. It supports interworking of disparate encapsulations at both ends and allows the connections to span multiple autonomous systems.

AToM requires that MPLS and Targeted Label Distribution Protocol (T-LDP) sessions are operational, which includes that the loopback addresses on PE routers use a /32 network mask and are not summarized in the core. Maximum transmission unit (MTU) considerations include the requirement of having identical MTU sizes on both sides, and guaranteeing sufficiently large MTU sizes in the core to avoid fragmentation. Fragmentation is not supported for Layer 2 VPNs.

Carrier Ethernet offers a range of enhancements for scalability (802.1ad and 802.1ah), high availability (MST, REP, and G.8032), and Operation, Administration, and Maintenance (OAM) (802.1ag and link OAM).

VPLS configuration uses virtual forwarding instances (VFIs) for bridging over the pseudowires (PWs) and can be combined with Cisco Ethernet Virtual Connection (EVC) capabilities for matching and manipulating frames.

https://t.me/learningnets

References

For additional information, refer to these resources:

- To learn more about implementing Layer 2 VPN on Cisco IOS XR platforms, refer to *Implementing Virtual Private LAN Services on Cisco IOS XR Software* at this URL: http://www.cisco.com/en/US/docs/routers/crs/software/crs_r3.9/lxvpn/configuration/guide/vc39vpls.pdf
- To learn more about configuring Layer 2 VPNs and Ethernet services on Cisco IOS XR platforms, refer to *Cisco ASR 9000 Series Aggregation Services Router L2VPN and Ethernet Services Configuration Guide* at this URL: http://www.cisco.com/en/US/docs/routers/asr9000/software/asr9k_r4.1/lxvpn/configuration/guide/lesc41.html

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets

Module Self-Check

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

Q1) To establish VPWS in your service provider environment, MPLS must be configured in your IP core network. (Source: Introducing Layer 2 VPNs)

A) true
B) false

Q2) To establish VPLS in your service provider environment, MPLS must be configured in your IP core network. (Source: Introducing Layer 2 VPNs)

A) true
B) false

Q3) A 1000-byte IPv6 packet is transported over EoMPLS. The control word is enabled. MPLS network is not configured for traffic engineering. The Layer 2 PDU of that packet in the core is ___ bytes. (Source: Introducing AToM)

Q4) A QinQ frame with outer VLAN tag 10 and inner VLAN tag 99 is transported over VPLS. The ingress PE has two interfaces with these settings:

- **encapsulation dot1q 10, rewrite ingress tag translate 1-to-1 dot1q 70 symmetric**
- **encapsulation dot1q 10 second 10-200, rewrite ingress tag pop 1 symmetric**

The frame transported over VPLS has this VLAN tagging: _____ (Source: Implementing VPLS)

https://t.me/learningnets

Module Self-Check Answer Key

Q1) B

Q2) A

Q3) 1012

Q4) 99 (only one VLAN tag)

 © 2012 Cisco Systems, Inc.

https://t.me/learningnets