

MPLS

Implementing Cisco MPLS

Volume 2

Version 2.1


Student Guide

Text Part Number: ILSG Production Services: 11.18.04

Copyright © 2004, Cisco Systems, Inc. All rights reserved.

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

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

 Copyright © 2004 Cisco Systems, Inc. All rights reserved. CCSP, the Cisco Square Bridge logo, Cisco Unity, Follow Me Browsing, FormShare, and StackWise are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, and iQuick Study are service marks of Cisco Systems, Inc.; and Aironet, ASIST, BPX, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Empowering the Internet Generation, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, GigaDrive, GigaStack, HomeLink, Internet Quotient, IOS, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, LightStream, Linksys, MeetingPlace, MGX, the Networkers logo, Networking Academy, Network Registrar, Packet, PIX, Post-Routing, Pre-Routing, ProConnect, RateMUX, Registrar, ScriptShare, SlideCast, SMARTnet, StrataView Plus, SwitchProbe, TeleRouter, The Fastest Way to Increase Your Internet Quotient, TransPath, and VCO are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website 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. (0406R)

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.

Table of Contents

Volume 2

<i>MPLS VPN Implementation</i>	5-1
Overview	5-1
Module Objectives	5-1
Using MPLS VPN Mechanisms of Cisco IOS Platforms	5-3
Overview	5-3
Objectives	5-3
What Is a Virtual Routing and Forwarding Table?	5-4
What Is the Need for Routing Protocol Contexts?	5-5
What Are VPN-Aware Routing Protocols?	5-6
How Are VRF Tables Used?	5-7
Propagating BGP Routes—Outbound	5-8
Example: BGP Route Propagation—Outbound	5-8
Example: BGP Route Propagation—Outbound	5-10
Propagating Non-BGP Routes—Outbound	5-11
Propagating Routes—Inbound	5-13
Summary	5-16
Configuring VRF Tables	5-17
Overview	5-17
Objectives	5-17
What Are the VRF Configuration Tasks?	5-18
Creating VRF Tables and Assigning RDs	5-19
ip vrf	5-19
Defaults	5-19
rd	5-20
Defaults	5-20
Specifying Export and Import RTs	5-21
route-target	5-21
Defaults	5-22
Assigning an Interface to a VRF Table	5-23
ip vrf forwarding	5-23
Defaults	5-23
Typical Configuration to Enable VRFs	5-24
Example: MPLS VPN Network	5-24
Summary	5-26
Configuring an MP-BGP Session Between PE Routers	5-27
Overview	5-27
Objectives	5-27
Configuring BGP Address Families	5-28
router bgp	5-29
Defaults	5-29
address-family	5-30
Enabling BGP Neighbors	5-31
Configuring MP-BGP	5-32
Configuring MP-IBGP	5-33
neighbor remote-as	5-34
Defaults	5-34
neighbor update-source	5-34
Defaults	5-34
neighbor activate	5-35
Defaults	5-35
neighbor next-hop-self	5-36
Defaults	5-36
Configuring MP-BGP Community Propagation	5-37
neighbor send-community	5-38
Defaults	5-38

Disabling IPv4 Route Exchange	5-40
Example: Disabling IPv4 Route Exchange	5-41
Summary	5-42
Configuring Small-Scale Routing Protocols Between PE and CE Routers	5-43
Overview	5-43
Objectives	5-43
Configuring PE-CE Routing Protocols	5-44
Selecting the VRF Routing Context for BGP	5-45
address-family ipv4	5-46
Defaults	5-46
Command Modes	5-46
Configuring Per-VRF Static Routes	5-47
ip route vrf	5-47
Configuring RIP PE-CE Routing	5-49
Configuring EIGRP PE-CE Routing	5-52
Summary	5-55
Monitoring MPLS VPN Operations	5-57
Overview	5-57
Objectives	5-57
Monitoring VRFs	5-58
show ip vrf	5-58
Defaults	5-58
Monitoring VRF Routing	5-62
show ip protocols vrf	5-62
show ip route vrf	5-63
show ip bgp vpnv4	5-64
show ip bgp vpnv4 vrf neighbors	5-67
Defaults	5-67
Usage Guidelines	5-67
Monitoring MP-BGP Sessions	5-68
show ip bgp neighbors	5-69
Example: Sample Output from show ip bgp neighbors Command	5-70
Monitoring an MP-BGP VPNv4 Table	5-72
show ip bgp vpnv4 vrf	5-73
Defaults	5-73
Usage Guidelines	5-73
show ip bgp vpnv4 rd <i>route-distinguisher</i>	5-74
Defaults	5-74
Usage Guidelines	5-75
Example: Configuring a Default RD for Two VRFs	5-75
Monitoring Per-VRF CEF and LFIB Structures	5-76
show ip cef vrf	5-77
Defaults	5-78
Usage Guidelines	5-78
show mpls forwarding vrf	5-79
Defaults	5-79
Usage Guidelines	5-80
Monitoring Labels Associated with VPNv4 Routes	5-81
Identifying Other MPLS VPN Monitoring Commands	5-82
Summary	5-83
Configuring OSPF as the Routing Protocol Between PE and CE Routers	5-85
Overview	5-85
Objectives	5-85
What Is the Enhanced OSPF Hierarchical Model?	5-86
Propagating OSPF Customer Routes	5-87
Implementing MPLS VPNs as an OSPF Superbackbone	5-90
Example: OSPF Superbackbone Implementation	5-95

Configuring OSPF PE-CE Routing	5-98
router ospf	5-99
Defaults	5-99
Using the OSPF Down Bit	5-101
Example: OSPF Down Bit	5-101
Example: OSPF Down Bit	5-103
Optimizing Packet Forwarding Across the MPLS VPN Backbone	5-104
Example: Optimizing of Packet Forwarding	5-104
Using the OSPF Tag Field	5-107
Example: OSPF Tag Field	5-107
Example: OSPF Tag Field—Routing Loop Prevention	5-110
What Is a Sham Link?	5-111
Example: Sham Link	5-111
Configuring a Sham Link	5-115
Defaults	5-116
Command Modes	5-116
Example: Sample Sham-Link Configuration	5-117
Summary	5-118
Configuring BGP as the Routing Protocol Between PE and CE Routers	5-119
Overview	5-119
Objectives	5-119
Configuring a Per-VRF BGP Routing Context	5-120
address-family ipv4	5-121
Defaults	5-121
Command Modes	5-121
Example: Configuring per-VRF BGP Routing Context	5-122
What Are the Reasons for Limiting the Number of Routes in a VRF?	5-123
Limiting the Number of Prefixes Received from a BGP Neighbor	5-124
neighbor maximum-prefix	5-124
Defaults	5-125
Limiting the Total Number of VRF Routes	5-126
maximum routes	5-127
Defaults	5-127
Example: Limiting the Total Number of VRF Routes	5-128
Identifying AS-Override Issues	5-129
neighbor as-override	5-132
Defaults	5-132
Example: AS-Override	5-133
Example: AS-Path Prepending	5-134
Identifying Allowas-in Issues	5-135
Example: Allowas-in	5-138
neighbor allowas-in	5-139
Defaults	5-139
Implementing SOO for Loop Prevention	5-140
set extcommunity	5-142
Defaults	5-143
neighbor route-map	5-143
ip vrf sitemap	5-144
Defaults	5-144
Summary	5-146
Troubleshooting MPLS VPNs	5-147
Overview	5-147
Objectives	5-147
Identifying Preliminary Steps in MPLS VPN Troubleshooting	5-148
Verifying the Routing Information Flow	5-149
Validating CE-to-PE Routing Information Flow	5-150
Validating PE-to-PE Routing Information Flow	5-151
Validating PE-to-CE Routing Information Flow	5-156
Identifying the Issues When Verifying the Data Flow	5-157

Validating CEF Status	5-158
show cef interface	5-159
Usage Guidelines	5-159
Validating the End-to-End LSP	5-162
Validating the LFIB Status	5-163
Summary	5-164
Module Summary	5-165
References	5-166
Module Self-Check	5-167
Module Self-Check Answer Key	5-176

Complex MPLS VPNs **6-1**

Overview	6-1
Module Objectives	6-1

Using Advanced VRF Import and Export Features **6-3**

Overview	6-3
Objectives	6-3
What Are Advanced VRF Features?	6-4
Configuring Selective VRF Import	6-5
import map	6-6
Defaults	6-6
Example: Configuring Selective VRF Import	6-7
Configuring Selective VRF Export	6-8
set extcommunity	6-9
Defaults	6-10
export map	6-10
Defaults	6-10
Example: Configuring Selective VRF Export	6-11
Summary	6-12

Introducing Overlapping VPNs **6-13**

Overview	6-13
Objectives	6-13
Who Are the Participants in Overlapping VPNs?	6-14
What Are Typical Overlapping VPN Usages?	6-15
Overlapping VPN Routing	6-16
Example: Overlapping VPN Routing	6-16
Overlapping VPN Data Flow	6-18
Configuring Overlapping VPNs	6-19
Example: Overlapping VPNs—Configuration Tasks	6-19
Example: Configuring Overlapping VPN VRFs	6-21
Summary	6-22

Introducing Central Services VPNs **6-23**

Overview	6-23
Objectives	6-23
What Are the Access Characteristics of a Central Services VPN?	6-24
What Are the Routing Characteristics of a Central Services VPN?	6-25
Example: Central Services VPN Routing	6-25
Identifying the Central Services VPN Data Flow Model	6-27
Configuring a Central Services VPN	6-28
Example: Configuring a Central Services VPN	6-30
Integrating a Central Services VPN with a Simple VPN	6-31
Identifying the RD Requirements When Integrating a Central Services and Simple VPN	6-33
Identifying the RT Requirements When Integrating Central Services and Simple VPN	6-34
Example: Configuring VRFs in a Central Services and Simple VPN	6-36
Summary	6-37

Introducing Managed CE Routers Service	6-39
Overview	6-39
Objectives	6-39
What Are the Requirements of Managed CE Routers?	6-40
What Are the VRF and RD Requirements?	6-41
Configuring Managed CE Routers	6-42
Example: Configuring VRFs	6-43
Summary	6-44
Introducing MPLS Managed Services	6-45
Overview	6-45
Objectives	6-45
What Are MPLS VPN Managed Services?	6-46
What Is Network Address Translation?	6-49
Example: Network Address Translation	6-52
Example: NAT Implementation with Multiple NAT Pools	6-53
What Is DHCP Relay?	6-54
Example: DHCP Relay—Corporate DHCP Server	6-56
Example: DHCP Relay—Shared DHCP Server	6-58
Example: DHCP Relay Configuration	6-59
What Are On-Demand Address Pools?	6-60
Example: On-Demand Address Pools	6-63
What Are HSRP and VRRP?	6-64
Example: HSRP and VRRP Today	6-66
What Are Multicast VPNs?	6-67
Example: Multicast VPNs—Default MDT	6-70
Example: Multicast VPNs—Data MDT	6-72
Enabling a VPN for Multicast Routing	6-74
Summary	6-76
Module Summary	6-77
References	6-77
Module Self-Check	6-78
Module Self-Check Answer Key	6-83
Internet Access from an MPLS VPN	7-1
Overview	7-1
Module Objectives	7-1
Introducing VPN Internet Access Topologies	7-3
Overview	7-3
Objectives	7-3
What Is Classical Internet Access for a VPN Customer?	7-4
What Are the Methods to Access the Internet from Every Customer Site?	7-8
What Is a Central Firewall Service?	7-11
What Is Wholesale Internet Access?	7-15
Summary	7-17
Introducing VPN Internet Access Implementation Methods	7-19
Overview	7-19
Objectives	7-19
Major Design Models	7-20
Internet Access Through Global Routing	7-21
Internet Access Through Separate Interfaces or Subinterfaces	7-22
Internet Access in VPNs	7-23
Summary	7-24

Separating Internet Access from VPN Services	7-25
Overview	7-25
Objectives	7-25
Internet Access Separated from VPNs	7-26
Implementing Separate Subinterfaces	7-27
Example: Internet Access Through a Dedicated Subinterface	7-28
Example: Internet Access Through a Dedicated Subinterface—Traffic Flow	7-29
Classical Internet Access for a VPN Customer	7-30
Accessing the Internet from Every Customer Site	7-31
Separate Internet Access	7-32
Summary	7-33
Implementing Internet Access as a Separate VPN	7-35
Overview	7-35
Objectives	7-35
What Is Internet Access as a Separate VPN?	7-36
Implementing Redundant Internet Access	7-38
Example: Redundant Internet Access	7-39
Implementing Classical Internet Access for a VPN Customer	7-40
Implementing Internet Access from Every Customer Site	7-41
Implementing Internet Access Through a Central Firewall Service	7-42
Implementing Wholesale Internet Access	7-43
Running an Internet Backbone in a VPN	7-44
Summary	7-45
Module Summary	7-46
References	7-46
Module Self-Check	7-47
Module Self-Check Answer Key	7-50

MPLS VPN Implementation

Overview

This module covers Multiprotocol Label Switching (MPLS) Virtual Private Network (VPN) implementation on Cisco IOS 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, and also advanced MPLS VPN-specific routing protocol features. This module continues with a description of MPLS VPN monitoring and debugging commands that are available on Cisco IOS platforms and concludes with a troubleshooting lesson describing failure scenarios, identifying symptoms, and providing remedial action.

Module Objectives

Upon completing this module, you will be able to configure, monitor, and troubleshoot VPN operations. This ability includes being able to meet these objectives:

- Describe the usage of VRF tables in an MPLS VPN environment
- Configure VRF tables
- Configure MP-BGP sessions between PE routers
- Configure small-scale routing protocols (static, RIP, and EIGRP) between CE and PE routers
- Monitor MPLS VPN operations
- Configure OSPF as the routing protocol between CE and PE routers
- Configure BGP as the routing protocol between CE and PE routers
- Troubleshoot VPN operations

Using MPLS VPN Mechanisms of Cisco IOS Platforms

Overview

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

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

Objectives

Upon completing this lesson, you will be able to describe the usage of VRF tables in an MPLS VPN environment. This ability includes being able to meet these objectives:

- Describe the characteristics of a VRF table
- Describe the need for routing protocol contexts
- Describe the characteristics of VPN-aware routing protocols
- Describe how VRF tables are used
- Describe the outbound BGP route propagation process in an MPLS VPN implementation
- Describe the outbound non-BGP route propagation process in an MPLS VPN implementation
- Describe the inbound route propagation process in an MPLS VPN implementation

What Is a Virtual Routing and Forwarding Table?

This topic describes the characteristics of a VRF table.

Virtual Routing and Forwarding Table

Cisco.com

- **A VRF is the routing and forwarding instance for a set of sites with identical connectivity requirements.**
- **Data structures associated with a VRF are as follows:**
 - IP routing table
 - CEF table
 - Set of rules and routing protocol parameters (routing protocol contexts)
 - List of interfaces that use the VRF
- **Other information associated with a VRF is as follows:**
 - Route distinguisher
 - Set of import and export route targets

© 2004 Cisco Systems, Inc. All rights reserved. MPLS v2.1—5-3

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

- The global IP routing table in Cisco IOS software
- A Cisco Express Forwarding (CEF) table identical in function to the global CEF forwarding table (Forwarding Information Base [FIB])
- Specifications for routing protocols running inside the VRF instance

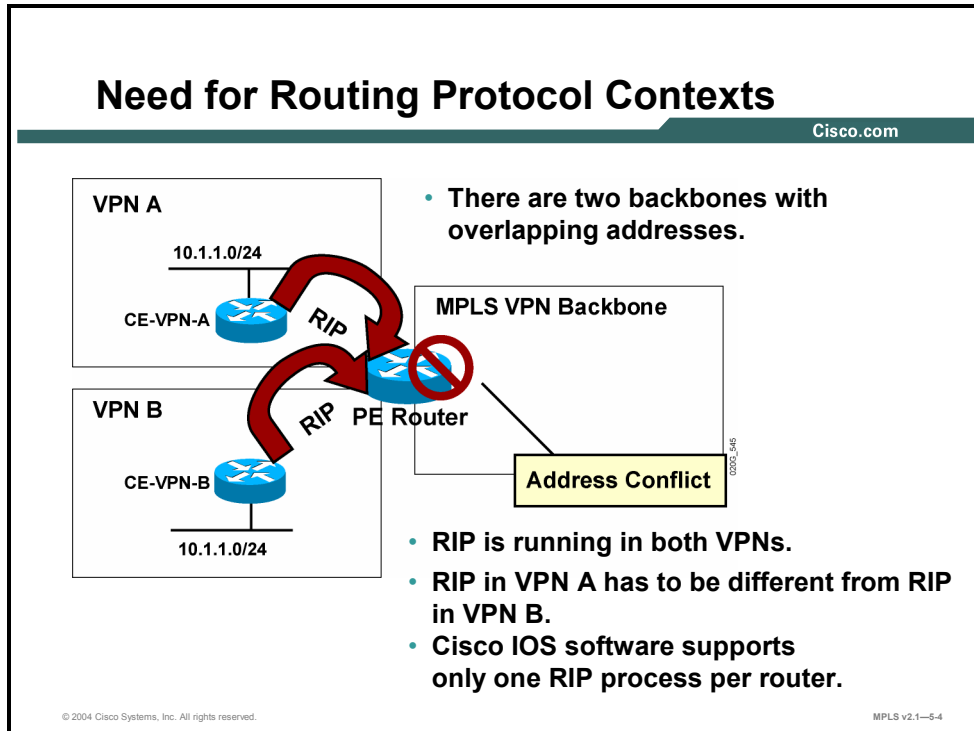
A VRF is thus a routing and forwarding instance that you can use for a single VPN site or for many sites connected to the same provider edge (PE) router *as long as these sites share exactly the same connectivity requirements.*

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

- The route distinguisher (RD), which is prepended (for example, RD + IP address) to all routes exported from the VRF into the global VPNv4—also called VPN IP version 4 (IPv4) Border Gateway Protocol (BGP) table
- A set of export route targets (RTs), which are attached to any route exported from the VRF
- A set of import RTs, which are used to select VPNv4 routes that are to be imported into the VRF

What Is the Need for Routing Protocol Contexts?

This topic describes the need for routing protocol contexts.



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

It is important to understand that for several important routing protocols, such as Routing Information Protocol (RIP), Enhanced Interior Gateway Routing Protocol (EIGRP), or BGP, Cisco IOS software supports only a single copy of the protocol running in the router. These protocols cannot be used directly between PE and customer edge (CE) routers in VPN environments because each VPN (or, more precisely, each VRF) needs a separate, isolated copy of the routing protocol to prevent undesired route leakage between VPNs. Furthermore, VPNs can use overlapping IP address spaces (for example, each VPN could use subnetworks of network 10.0.0.0), which would also lead to routing confusions if all VPNs shared the same copy of the routing protocol.

What Are VPN-Aware Routing Protocols?

This topic describes the characteristics of VPN-aware routing protocols.

VPN-Aware Routing Protocols

Cisco.com

Routing context = routing protocol run in one VRF:

- **Supported by VPN-aware routing protocols:**
 - External BGP (EBGP), EIGRP, OSPF, RIP version 2 (RIPv2), static routes
- **Implemented as several instances of a single routing process (EBGP, RIPv2) or as several routing processes (OSPF)**
- **Independent per-instance router variables for each instance**

© 2004 Cisco Systems, Inc. All rights reserved.MPLS v2.1—5-5

“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 (OSPF), similar to traditional Cisco IOS software implementation, or as separate isolated “instances” of the same routing protocol.

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

How Are VRF Tables Used?

This topic describes how VRF tables are used in an MPLS VPN implementation.

VRF Table

Cisco.com

- **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**

© 2004 Cisco Systems, Inc. All rights reserved.MPLS v2.1—5-6

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


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

Note The requirement to support CEF switching on inbound VRF interfaces prevents certain media or encapsulation types from being used for VPN connectivity. More notable examples in mainstream Cisco IOS Release 12.1 include dialer interfaces, ISDN interfaces, and Switched Multimegabit Data Service (SMDS) interfaces. Some restrictions are already lifted in Cisco IOS Release 12.1 T. Refer to the release notes of the Cisco IOS platform that you are using for details about the interfaces and media types supporting CEF switching.

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

Propagating BGP Routes—Outbound

This topic describes the outbound BGP route propagation process in an MPLS VPN implementation.




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

Example: BGP Route Propagation—Outbound

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

BGP Route Propagation—Outbound (Cont.)

Cisco.com



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

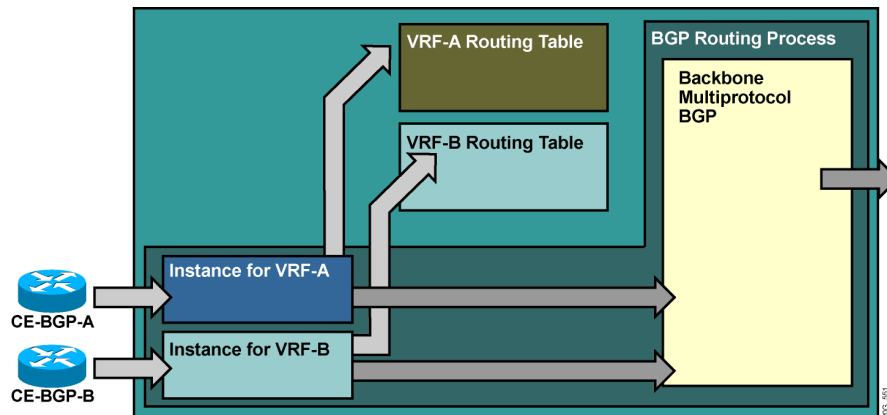
© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-8

The BGP-speaking CE routers announce their networks via EBGP sessions to the PE router. The CE routers are BGP neighbors of the PE router. The PE router associates each BGP neighbor relationship with individual VRFs that enable the various instances of the BGP routing process to put the received routing updates into the proper per-VRF routing table.

BGP Route Propagation—Outbound (Cont.)

Cisco.com



- The route distinguisher is prepended during the route export to the BGP routes from the VRF instance of BGP process to convert them into VPNv4 prefixes. Route targets are attached to these prefixes.
- VPNv4 prefixes are propagated to other PE routers.

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-9

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

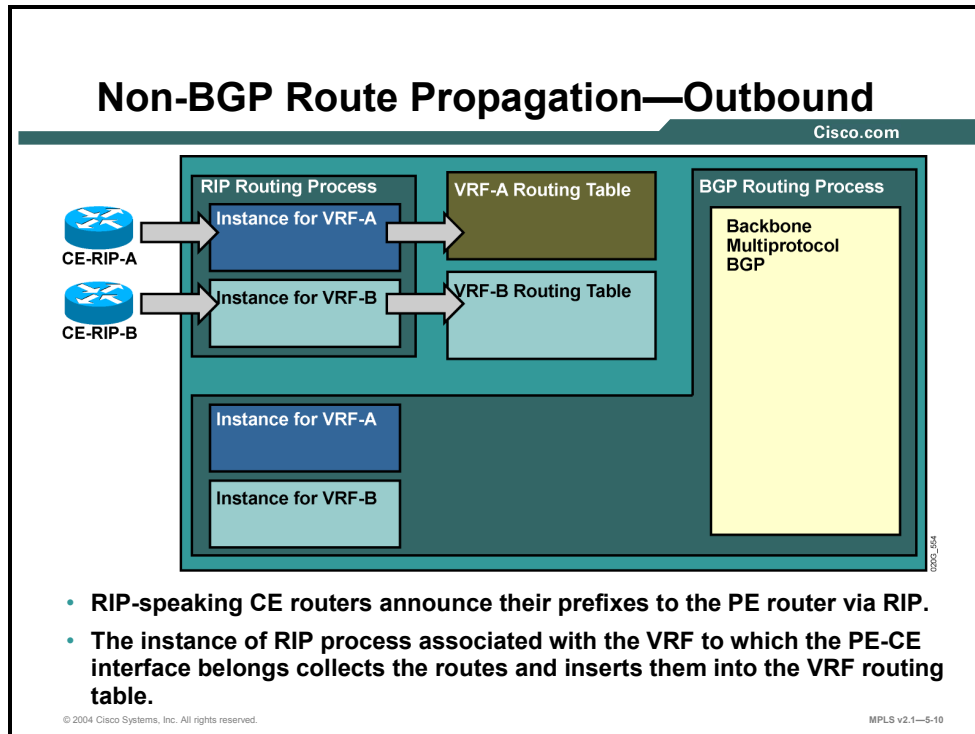
The routes illustrated here are being copied into the MP-BGP table for further propagation to other PE routers.

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

Note The difference between the per-VRF BGP table and the global MP-BGP table holding VPNv4 routes is displayed only to illustrate the steps in the route propagation process. In reality, there is no separate per-VRF BGP table in Cisco IOS software.

Propagating Non-BGP Routes—Outbound


This topic describes the outbound non-BGP route propagation process in an MPLS VPN implementation.



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

Non-BGP Route Propagation—Outbound (Cont.)

Cisco.com



- 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.**

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-11


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

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

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

Propagating Routes—Inbound

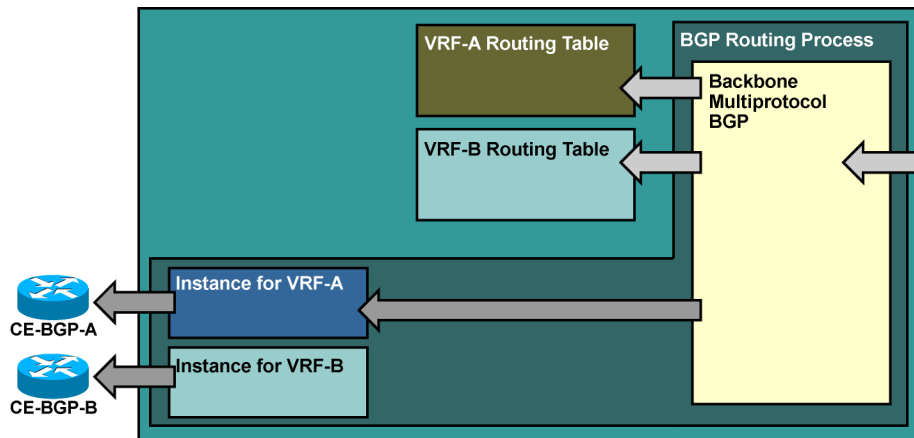
This topic describes the inbound route propagation process in an MPLS VPN implementation.



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

Route Propagation—Inbound (Cont.)

Cisco.com



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

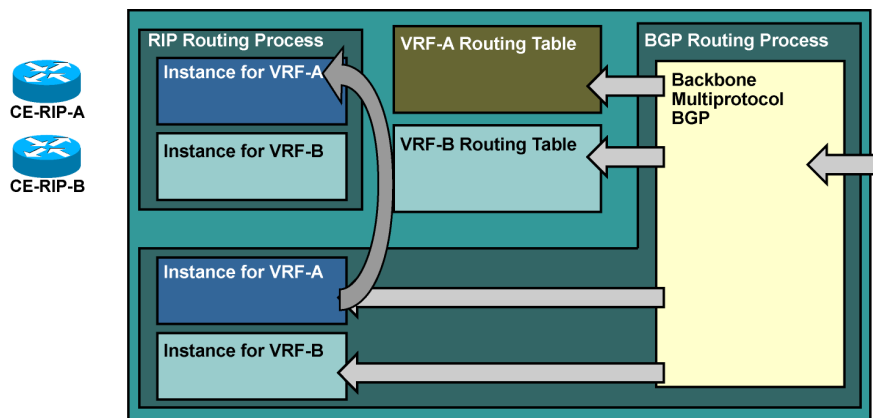
© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-13

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

Route Propagation—Inbound (Cont.)

Cisco.com



- MP-IBGP routes imported into a VRF are redistributed into the instance of RIP configured for that VRF.
- **Redistribution between BGP and RIP has to be configured for end-to-end RIP routing between CE routers.**


© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-14

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

Route Propagation—Inbound (Cont.)

Cisco.com



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

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-15

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

Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- **A VRF table is a routing and forwarding instance that associates additional attributes such as RD, import RT, and export RT to routing entries.**
- **Routing contexts allow multiple copies of routing protocols to run concurrently as separate VRF instances to prevent undesired route leakage between VPNs.**
- **VPN-aware routing protocols allow separation of routing tables either as separate routing processes (OSPF) or separate isolated instances of the same protocol (BGP, EIGRP, RIPv2).**
- **A VRF table is used to logically separate routing information from different VPNs.**

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-16

Summary (Cont.)

Cisco.com

- **Outbound BGP route propagation starts with CE BGP updates. Because the protocol source is BGP, MP-BGP can directly prepend RD and RT to the respective inbound instances of CE BGP updates.**
- **Outbound non-BGP route propagation starts with CE protocols other than BGP. Therefore, an additional step of redistribution is required before prepending RD and RT.**
- **Inbound route propagation filters routes based on RT into respective instances of VRF. Non-BGP speaking CE routers require redistribution.**

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-17

Configuring VRF Tables

Overview

This lesson explains how to configure VRF tables, listing the configuration tasks, syntax, and definitions of commands used for the creation of VRFs. The lesson also 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 will provide you with the information that will enable you to succeed at such tasks.

Objectives

Upon completing this lesson, you will be able to describe how to configure VRF tables. This ability includes being able to meet these objectives:

- Identify the tasks that are required to configure a VRF table
- Create a VRF table and assign RDs
- Specify export and import RTs
- Assign an interface to a VRF table
- Describe a typical Cisco IOS configuration that enables VRFs

What Are the VRF Configuration Tasks?

This topic identifies the tasks required to configure a VRF table.

VRF Configuration Tasks

Cisco.com

VRF configuration tasks:

- Create a VRF table.
- Assign RD to the VRF.
- Specify export and import route targets.
- Assign interfaces to VRFs.

© 2004 Cisco Systems, Inc. All rights reserved.MPLS v2.1—5-3

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

- Create a new VRF table.
- Assign a unique RD to the VRF.

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

- Specify import and export RTs for the VRF.

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

- Assign interfaces to VRFs.

Creating VRF Tables and Assigning RDs

This topic describes how to create a VRF table and assign RDs.

Creating VRF Tables and Assigning RDs

Cisco.com

Router (config) #

`ip vrf name`

- This command creates a new VRF or enters configuration of an existing VRF.
- VRF names are case-sensitive.
- VRF is not operational unless you configure RD.
- VRF names have only local significance.

Router (config-vrf) #

`rd route-distinguisher`

- This command assigns a route distinguisher to a VRF.
- You can use ASN:nn or A.B.C.D:nn format for RD.
- Each VRF in a PE router has to have a unique RD.

© 2004 Cisco Systems, Inc. All rights reserved. MPLS v2.1—5-4

ip vrf

To configure a VRF routing table, use the **ip vrf** command in global configuration mode. To remove a VRF routing table, use the **no** form of this command.

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

This table describes the parameters for the **ip vrf** command.

Syntax Description

Parameter	Description
<i>vrf-name</i>	Name assigned to a VRF.

Defaults

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

rd

To create routing and forwarding tables for a VRF, use the **rd** command in VRF configuration submode: **rd** *route-distinguisher*.

This table describes the parameters for the **rd** command.

Syntax Description

Parameter	Description
<i>route-distinguisher</i>	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)

Defaults

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

Note	Once a VRF has been defined using the ip vrf command and a RD has been assigned using the rd command, the VRF is operational. At this point, any locally active interface will appear in the routing display of the VRF table.
-------------	--

Specifying Export and Import RTs

This topic describes how to specify export and import RTs.

Specifying Export and Import RTs

Cisco.com

```
Router (config-vrf) #  
route-target export RT
```

- Specifies an RT to be attached to every route exported from this VRF to MP-BGP.
- Allows specification of many export RTs—all to be attached to every exported route.

```
Router (config-vrf) #  
route-target import RT
```

- Specifies an RT to be used as an import filter—only routes matching the RT are imported into the VRF.
- Allows specification of many import RTs—any route where at least one RT attached to the route matches any import RT is imported into the VRF.

Because of implementation issues, at least one export route target must also be an import route target of the same VRF in Cisco IOS Release 12.0 T.

© 2004 Cisco Systems, Inc. All rights reserved. MPLS v2.1—5-5

route-target

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

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

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

Syntax Description

Parameter	Description
import	VPNv4 routes that contain an extended community value that matches the route-target-ext-community field that will be imported into the VRF.
export	The value in the route-target-ext-community field that will be inserted into the extended community for routes exported from the VRF to VPNv4.
both	Sets the value used by both the import and export process to the value indicated in the route-target-ext-community field.
<i>route-target-ext-community</i>	The RT extended community attribute for the VRF.

Similar to RDs, the 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)

Defaults

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

Specifying Export and Import RTs (Cont.)

Cisco.com

```
Router(config-vrf)#  
route-target both RT
```

- In cases where the export RT matches the import RT, use this form of the route-target command.

Sample router configuration for simple customer VPN:

```
ip vrf Customer_ABC  
rd 12703:15  
route-target export 12703:15  
route-target import 12703:15
```

© 2004 Cisco Systems, Inc. All rights reserved. MPLS v2.1—5-6

Whenever an RT is both an import and an export RT for a VRF, you can use the **route-target both** command to simplify the configuration. For example, the “route-target” configuration lines in the sample router configuration in the figure could be reduced to a single command: **route-target both 12703:15**.

Assigning an Interface to a VRF Table

This topic describes how to assign an interface to a VRF table.

Assigning an Interface to VRF Table

Cisco.com

```
Router(config-if)#  
ip vrf forwarding vrf-name
```

- This command associates an interface with the specified VRF.
- The existing IP address is removed from the interface when interface is put into VRF—the IP address must be reconfigured.
- CEF switching must be enabled on the interface.

Sample router configuration:

```
ip cef  
!  
interface serial 0/0  
ip vrf forwarding Customer_ABC  
ip address 10.0.0.1 255.255.255.252
```

© 2004 Cisco Systems, Inc. All rights reserved.MPLS v2.1—5-7

ip vrf forwarding

To associate a VRF with an interface or subinterface, use the **ip vrf forwarding** command in interface configuration mode. To disassociate a VRF, use the **no** form of this command.

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

This table describes the parameters for the **ip vrf forwarding** command.

Syntax Description


Parameter	Description
<i>vrf-name</i>	Name assigned to a VRF.

Defaults

The default for an interface is the global routing table.

Typical Configuration to Enable VRFs

This topic describes a typical Cisco IOS configuration that enables VRFs.



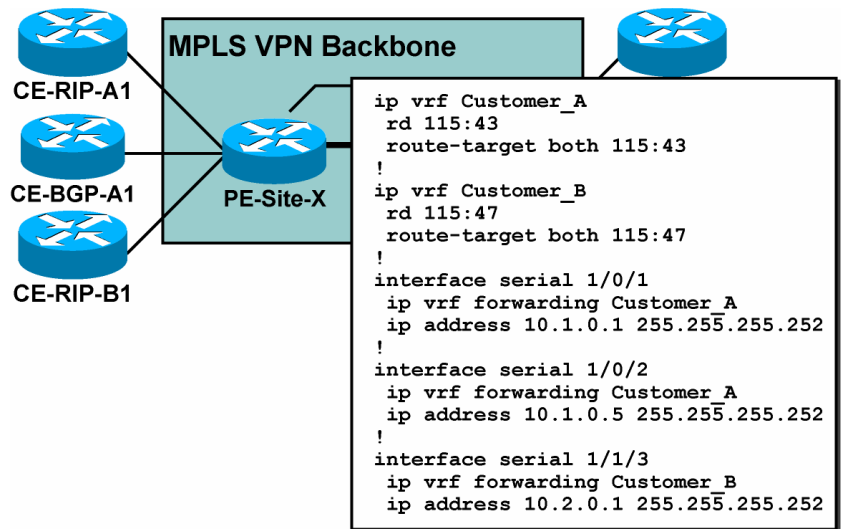
To illustrate the use of MPLS VPN configuration commands, you can look at a configuration of the PE router in a sample network.

Example: MPLS VPN Network

The figure illustrates a configuration of the PE router in a sample network with two VPN customers. Customer A (with four sites) is using BGP and RIP as the provider edge-customer edge (PE-CE) routing protocol, and customer B (with two sites) is using only RIP. Both customers use private IP address space (subnetworks of network 10.0.0.0).

MPLS VPN Network Example (Cont.)

Cisco.com



© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-9

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

- Step 1** Configure VRFs for customer A and customer B.
- Step 2** Assign RDs and RTs to the VRFs. Only one RD per customer is used on all PE routers in the MPLS VPN backbone, because these customers require only simple VPN connectivity. To simplify the configuration and troubleshooting process, the RTs are made equal to the RDs.
- Step 3** Assign PE-CE interfaces to individual VRFs.

Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- **There are four VRF configuration tasks: create a VRF table, assign RD to the VRF, specify export and import route targets, and assign interfaces to VRFs.**
- **Use the ip vrf command to configure a VRF routing table.**
- **Use the route-target command to import and export between the VRF and MP-BGP with respective RT values.**
- **Use the ip vrf forwarding command to associate a VRF with an interface or subinterface.**
- **In a typical Cisco IOS configuration, you will:**
 - **Configure VRFs for the customers.**
 - **Assign RDs and RTs to the VRFs.**
 - **Assign PE-CE interfaces to individual VRFs.**

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-10

Configuring an MP-BGP Session Between PE Routers

Overview

This lesson explains the BGP process in an MPLS VPN-enabled router, listing the configuration tasks, steps, syntax, and descriptions. The lesson also discusses BGP community propagation and provides an MP-IBGP configuration example.

Most of the configuration in an MPLS VPN depends on how the PE routers are configured. Having a good grasp of exactly what is being configured and why will help greatly to ensure that your MPLS VPN network operates as smoothly as possible.

Objectives

Upon completing this lesson, you will be able to describe how to configure MP-BGP in an MPLS VPN backbone. This ability includes being able to meet these objectives:

- Configure BGP address families
- Describe the requirements for enabling BGP neighbors in an MPLS VPN environment
- Identify the process steps involved in configuring MP-BGP in an MPLS VPN environment
- Configure MP-IBGP in an MPLS VPN environment
- Configure MP-BGP community propagation in an MPLS VPN environment
- Disable IPv4 route exchange in an MPLS VPN environment

Configuring BGP Address Families

This topic describes how to configure BGP address families.

Configuring BGP Address Families

Cisco.com

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

© 2004 Cisco Systems, Inc. All rights reserved.MPLS v2.1—5-3

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

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

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

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

Configuring BGP Address Families (Cont.)

Cisco.com

Router(config)#

```
router bgp as-number
```

- **Selects global BGP routing process.**

Router(config-router)#

```
address-family vpnv4
```

- **Selects configuration of VPNv4 prefix exchanges under MP-BGP sessions.**

Router(config-router)#

```
address-family ipv4 vrf vrf-name
```

- **Selects configuration of per-VRF PE-CE EBGP parameters.**

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-4

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

- 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.

router bgp

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

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

This table describes the **router bgp** command.

Syntax Description

Parameter	Description
<i>as-number</i>	Displays the number of an AS that identifies the router to other BGP routers and tags the routing information passed along.

Defaults

No BGP routing process is enabled by default.

address-family

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

- VPNv4 unicast:
 - **address-family vpnv4 [unicast]**
 - **no address-family vpnv4 [unicast]**
- IPv4 unicast:
 - **address-family ipv4 [unicast]**
 - **no address-family ipv4 [unicast]**
- IPv4 unicast with CE router:
 - **address-family ipv4 [unicast] vrf *vrf-name***
 - **no address-family ipv4 [unicast] vrf *vrf-name***

This table describes the **address-family** command.

Syntax Description

Parameter	Description
ipv4	Configures sessions that carry standard IPv4 address prefixes.
vpnv4	Configures sessions that carry customer VPNv4 prefixes, each of which has been made globally unique by adding an 8-byte (B) RD.
unicast	(Optional) Specifies unicast prefixes.
vrf <i>vrf-name</i>	Specifies the name of a VPN VRF to associate with submode commands.

Enabling BGP Neighbors

This topic describes the requirements for enabling BGP neighbors in an MPLS VPN environment.

BGP Neighbors

Cisco.com

- **MP-BGP neighbors are configured under the BGP routing process:**
 - These neighbors need to be activated for each global address family that they support.
 - Per-address-family parameters can be configured for these neighbors.
- **VRF-specific EBGP neighbors are configured under corresponding address families.**

© 2004 Cisco Systems, Inc. All rights reserved.MPLS v2.1—5-5

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 only have to be *activated* for individual address families.
- Per-VRF BGP neighbors (the CE routers), which are configured and activated within the **address-family ipv4 vrf vrf-name** command

Configuring MP-BGP

This topic identifies the process steps involved in configuring MP-BGP in an MPLS VPN environment.

Configuring MP-BGP

Cisco.com

MPLS VPN MP-BGP configuration steps:

- **Configure MP-BGP neighbor under BGP routing process.**
- **Configure BGP address family VPNv4.**
- **Activate configured BGP neighbor for VPNv4 route exchange.**
- **Specify additional parameters for VPNv4 route exchange (filters, next hops, and so on).**

© 2004 Cisco Systems, Inc. All rights reserved.MPLS v2.1—5-6

Configure BGP connectivity between two PE routers in these four steps:

- Step 1** Configure the remote PE router as a global BGP neighbor in BGP router configuration mode.
- Step 2** Define the parameters that affect all BGP route exchange (for example, source address for the TCP session) on the global BGP neighbor.
- Step 3** Select the VPNv4 address family and activate the BGP neighbor for VPNv4 route exchange.
- Step 4** Configure additional VPNv4-specific BGP parameters (filters, next-hop processing, route maps) within the VPNv4 address family.

Note IPv4-specific BGP parameters are still configured under the BGP router configuration mode; there is no special IPv4 address family.

Configuring MP-IBGP

This topic describes how to configure MP-IBGP in an MPLS VPN environment.

Configuring MP-IBGP

Cisco.com

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

- All MP-BGP neighbors have to be configured under 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.

© 2004 Cisco Systems, Inc. All rights reserved.MPLS v2.1—5-7

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

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

neighbor remote-as

To add an entry to the BGP neighbor table, use the **neighbor remote-as** command in router configuration mode. To remove an entry from the table, use the **no** form of this command.

- **neighbor** {*ip-address* | *peer-group-name*} **remote-as** *as-number*
- **no neighbor** {*ip-address* | *peer-group-name*} **remote-as** *as-number*

This table describes the **neighbor remote-as** command.

Syntax Description

Parameter	Description
<i>ip-address</i>	Neighbor IP address.
<i>peer-group-name</i>	Name of BGP peer group.
<i>as-number</i>	AS to which the neighbor belongs.

Defaults

There are no BGP neighbor peers.

neighbor update-source

To have the Cisco IOS software allow internal BGP sessions to use any operational interface for TCP connections, use the **neighbor update-source** command in router configuration mode. To restore the interface assignment to the closest interface, which is called the “best local address,” use the **no** form of this command.

- **neighbor** {*ip-address* | *peer-group-name*} **update-source** *interface-type*
- **no neighbor** {*ip-address* | *peer-group-name*} **update-source** *interface-type*

This table describes the **neighbor update-source** command.

Syntax Description

Parameter	Description
<i>ip-address</i>	IP address of the BGP-speaking neighbor.
<i>peer-group-name</i>	Name of BGP peer group.
<i>interface-type</i>	Loopback interface.

Defaults

The default is the best local address.

Configuring MP-IBGP (Cont.)

Cisco.com

```
Router(config-router-af)#
```

```
neighbor ip-address activate
```

- The BGP neighbor defined under BGP router configuration has to be activated for VPNv4 route exchange.

```
Router(config-router-af)#
```

```
neighbor ip-address next-hop-self
```

- The next-hop-self keyword can be configured on the MP-IBGP session for MPLS VPN configuration if EBGP is being run with a CE neighbor.

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-8

After you define the remote PE router as a global BGP neighbor, you must activate it for VPNv4 route exchange.

neighbor activate

To enable the exchange of information with a BGP neighboring router, use the **neighbor activate** command in router configuration mode. To disable the exchange of an address with a neighboring router, use the **no** form of this command.

- **neighbor** {*ip-address* | *peer-group-name*} **activate**
- **no neighbor** {*ip-address* | *peer-group-name*} **activate**

This table describes the **neighbor activate** command.

Syntax Description

Parameter	Description
<i>ip-address</i>	IP address of the neighboring router.
<i>peer-group-name</i>	Name of BGP peer group.

Defaults

The exchange of addresses with neighbors is enabled by default for the IPv4 address family. For all other address families, address exchange is disabled by default. You can explicitly activate the default command by using the appropriate address family submenu.

neighbor next-hop-self

To disable next-hop processing of BGP updates on the router, use the **neighbor next-hop-self** command in router configuration mode. To disable this feature, use the **no** form of this command.

- **neighbor** {*ip-address* | *peer-group-name*} **next-hop-self**
- **no neighbor** {*ip-address* | *peer-group-name*} **next-hop-self**

This table describes the **neighbor next-hop-self** command.

Syntax Description

Parameter	Description
<i>ip-address</i>	IP address of the BGP-speaking neighbor.
<i>peer-group-name</i>	Name of BGP peer group.

Defaults

Default is disabled.

Configuring MP-BGP Community Propagation

This topic describes how to configure MP-BGP community propagation in an MPLS VPN environment.

MP-BGP Community Propagation

Cisco.com

```
Router (config-router-af) #  
neighbor ip-address send-community [extended | both]
```

- This command configures propagation of standard and extended BGP communities attached to VPNv4 prefixes.
- Default value: only extended communities are sent.
- Usage guidelines:
 - Extended BGP communities attached to VPNv4 prefixes **have to be exchanged** between MP-BGP neighbors for proper MPLS VPN operation.
 - To propagate standard BGP communities between MP-BGP neighbors, use the **both** option.

© 2004 Cisco Systems, Inc. All rights reserved. MPLS v2.1—5-9

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

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

neighbor send-community

To specify that BGP community attributes that are attached to a BGP route should be sent to a BGP neighbor, use the **neighbor send-community** command in router configuration mode. To remove the entry, use the **no** form of this command.

- **neighbor** {*ip-address* | *peer-group-name*} **send-community** [**extended** | **both**]
- **no neighbor** {*ip-address* | *peer-group-name*} **send-community**

This table describes the **neighbor send-community** command.

Syntax Description


Parameter	Description
<i>ip-address</i>	Neighbor IP address.
<i>peer-group-name</i>	Name of BGP peer group.

Defaults

BGP communities are not propagated to any neighbor.

MP-BGP BGP Community Propagation (Cont.)

Cisco.com



© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-10

The configuration example provided in the “Configuring VRF Tables” lesson continues here with configuration of MP-IBGP sessions on the PE router. This table describes the steps that you need to perform.

Configuration of MP-IBGP sessions

Step	Action
1	Define a loopback interface that will serve as the BGP next hop for VPNv4 routes and as the source address for the IBGP session.
2	Configure the remote PE router as the global BGP neighbor.
3	Specify the source address for the TCP session.
4	Select the VPNv4 address family.
5	Activate the remote PE router for VPNv4 route exchange.
6	Disable next-hop processing for VPNv4 route exchange. This action guarantees that the loopback 0 interface will always be the BGP next hop for VPNv4 routes propagated by this router to its MP-IBGP neighbors.
7	Configure propagation of standard and extended communities.

Disabling IPv4 Route Exchange

This topic describes how to disable IPv4 route exchange in an MPLS VPN environment.

Disabling IPv4 Route Exchange

Cisco.com

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

- **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.**

© 2004 Cisco Systems, Inc. All rights reserved.MPLS v2.1—5-11

The BGP configuration discussed so far is appropriate for scenarios where the PE routers provide Internet and VPN connectivity. If the PE routers provide only VPN connectivity, they do not need Internet routing and the IPv4 route exchange should be disabled. 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.

Disabling IPv4 Route Exchange (Cont.)

Cisco.com

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

```
router bgp 12703
no bgp default ipv4-unicast
neighbor 172.16.32.14 remote-as 12703
neighbor 172.16.32.15 remote-as 12703
neighbor 172.16.32.27 remote-as 12703

! Activate IPv4 route exchange
neighbor 172.16.32.14 activate
neighbor 172.16.32.27 activate

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

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-12

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

Example: Disabling IPv4 Route Exchange

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

Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- **Use the address-family command to select the routing context that you want to configure. Use the router bgp command in global configuration mode to configure the BGP routing process.**
- **VRF-specific EBGP neighbors are configured under corresponding address families.**
- **To configure MPLS VPN MP-BGP, you need to:**
 - **Configure MP-BGP neighbors.**
 - **Configure MP-BGP address family to start VPNv4 routing.**
 - **Activate configured MP-BGP neighbors.**
 - **Specify additional parameters for VPNv4 route exchange.**
- **The commands used to configure MP-BGP are: neighbor remote-as, neighbor update-source, neighbor activate, and neighbor next-hop-self.**
- **Use the neighbor send-community command to support standard and extended communities.**
- **There are two ways to disable IPv4 route exchange: use the no neighbor activate command or the no bgp default ipv4-unicast command.**

© 2004 Cisco Systems, Inc. All rights reserved.MPLS v2.1—5-13

Configuring Small-Scale Routing Protocols Between PE and CE Routers

Overview

This lesson explains the PE-CE routing protocol configuration steps and the various routing protocols that you can run between PE and CE routers. These protocols include RIP, EIGRP, and static routes.

It is important to understand not only what you can configure between PE and CE routers when you are setting up MPLS VPNs, but also how to accomplish the configuration successfully. This lesson looks at the configuration parameters that you need to configure an MPLS VPN PE-CE routing exchange.

Objectives

Upon completing this lesson, you will be able to describe how to configure small scale routing protocols between PE and CE routers. This ability includes being able to meet these objectives:

- Identify the requirements for configuring PE-CE routing protocols
- Select the VRF routing context for BGP
- Configure per-VRF static routes
- Configure a RIP PE-CE routing session
- Configure an EIGRP PE-CE routing session

Configuring PE-CE Routing Protocols

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

PE-CE Routing Protocols

Cisco.com

- PE-CE routing protocols are configured for individual VRFs.
- 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.
- The overall number of routing processes per router is limited to 32, of which **only 28** are available for VRF assignment.

© 2004 Cisco Systems, Inc. All rights reserved.MPLS v2.1–5-3

After you configure VRFs and establish MP-IBGP connectivity between PE routers, you have to configure routing protocols between the PE router and the attached CE routers. The PE-CE routing protocols need to be configured for individual VRFs. Sites in the same VPN but in different VRFs cannot share the same PE-CE routing protocol.

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

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

- Per-VRF routing protocols can be configured as individual address families belonging to the same routing process (similar to what you have already seen for BGP).
- Per-VRF routing protocols can be configured as separate routing processes. This option is used for more complex routing protocols that need to maintain a separate topology database for each VRF (for example, OSPF).

Note Current Cisco IOS software implementation limits the overall number of routing protocols in a router to 32. Two routing methods are predefined (static and connected), and two routing protocols are needed for proper MPLS VPN backbone operation—BGP and backbone Interior Gateway Protocol (IGP). The number of PE-CE routing processes is therefore limited to 28.

Selecting the VRF Routing Context for BGP

This topic describes how to select the VRF routing context for BGP.

Configuring the VRF Routing Context Within BGP

Cisco.com

```
Router (config) #  
router bgp as-number  
  address-family ipv4 vrf vrf-name  
    ... Non-BGP redistribution ...
```

- **Select the per-VRF BGP context with the address-family command.**
- **Configure CE EBGP neighbors in VRF context, not in global BGP configuration.**
- **All non-BGP per-VRF routes have to be redistributed into per-VRF BGP context to be propagated by MP-BGP to other PE routers.**

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-4

Select the VRF routing context with the **address-family ipv4 vrf *vrf-name*** command in the RIP and BGP routing processes. All per-VRF routing protocol parameters (network numbers, passive interfaces, neighbors, filters, and so on) are configured under this address family.

Note Common parameters defined in router configuration mode are inherited by all address families defined for this routing process and can be overridden for each individual address family.

address-family ipv4

To enter address family configuration mode for configuring routing sessions, such as BGP, that use standard IPv4 address prefixes, use the **address-family ipv4** command in router configuration mode. To disable address family configuration mode, use the **no** form of this command.

- **address-family ipv4** [**multicast** | **unicast** | **vrf** *vrf-name*]
- **no address-family ipv4** [**multicast** | **unicast** | **vrf** *vrf-name*]

This table describes the **address-family ipv4** command.

Syntax Description

Parameter	Description
multicast	(Optional) Specifies IPv4 multicast address prefixes.
unicast	(Optional) Specifies IPv4 unicast address prefixes.
vrf <i>vrf-name</i>	(Optional) Specifies the name of the VRF instance to associate with subsequent IPv4 address family configuration mode commands.

Defaults

IPv4 address prefixes are not enabled. Unicast address prefixes are the default when IPv4 address prefixes are configured.

Command Modes

This command is used in router configuration mode.

Configuring Per-VRF Static Routes

This topic describes how to configure per-VRF static routes.

Configuring Per-VRF Static Routes

Cisco.com

```
Router(config)#  
ip route vrf name static route parameters
```

- This command configures per-VRF static routes.
- The route is entered in the VRF table.
- **You must always specify the outgoing interface, even if you specify the next hop.**

Sample router configuration:

```
ip route vrf Customer_ABC 10.0.0.0 255.0.0.0 10.250.0.2 serial 0/0  
!  
router bgp 12703  
 address-family ipv4 vrf Customer_ABC  
  redistribute static
```

© 2004 Cisco Systems, Inc. All rights reserved. MPLS v2.1—5-5

ip route vrf

To establish static routes for a VRF, use the **ip route vrf** command in global configuration mode. To disable static routes, use the **no** form of this command.

- **ip route vrf** *vrf-name prefix mask [next-hop-address] [interface {interface-number}] [global] [distance] [permanent] [tag tag]*
- **no ip route vrf** *vrf-name prefix mask [next-hop-address] [interface {interface-number}] [global] [distance] [permanent] [tag tag]*

This table describes the **ip route vrf** command.

Syntax Description

Parameter	Description
<i>vrf-name</i>	Name of the VRF for the static route.
<i>prefix</i>	IP route prefix for the destination, in dotted decimal notation.
<i>mask</i>	Prefix mask for the destination, in dotted decimal notation.
<i>next-hop-address</i>	(Optional) IP address of the next hop (the forwarding router that can be used to reach that network).
<i>interface</i>	Type of network interface to use: ATM, Ethernet, loopback, Packet over SONET (POS), or null.
<i>interface-number</i>	Number identifying the network interface to use.
global	(Optional) Specifies that the given next-hop address be in the non-VRF routing table.
<i>distance</i>	(Optional) An administrative distance for this route.
permanent	(Optional) Specifies that this route will not be removed, even if the interface shuts down.
tag <i>tag</i>	(Optional) Label (tag) value that can be used for controlling redistribution of routes through route maps.

Configuring RIP PE-CE Routing

This topic describes how to configure a RIP PE-CE routing session.

Configuring RIP PE-CE Routing

Cisco.com

- 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.

© 2004 Cisco Systems, Inc. All rights reserved.MPLS v2.1—5-6

Configuring RIP as the PE-CE routing protocol is even easier than configuring BGP. Start the configuration of individual routing context with the **address-family ipv4 vrf vrf-name** command in router configuration mode. You can enter all standard RIP parameters in the per-VRF routing context. Global RIP parameters entered in the scope of RIP router configuration are inherited by each routing context and can be overwritten if needed in each routing context.

Note Only RIP version 2 (RIPv2) is supported as the PE-CE routing protocol. It is good configuration practice to configure the RIP version as a global RIP parameter using the **version 2** command in router configuration mode.

Configuring RIP PE-CE Routing: RIP Metric Propagation

Cisco.com

Router (config) #

```
router rip
  address-family ipv4 vrf vrf-name
    redistribute bgp as-number metric transparent
```

- BGP routes must be redistributed back into RIP.
- The RIP hop count has to be manually set for routes redistributed into RIP.
- For end-to-end RIP networks, the following applies:
 - On the sending end, the RIP hop count is copied into the BGP multi-exit discriminator attribute (default BGP behavior).
 - On the receiving end, the metric transparent option copies the BGP MED into the RIP hop count, resulting in a consistent end-to-end RIP hop count.
- When you are using RIP with other protocols, the metric must be manually set.

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-7

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


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

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

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

Configuring RIP PE-CE Routing: Example

Cisco.com



```
router rip
version 2
address-family ipv4 vrf Customer-RIP-A1
network 10.0.0.0
redistribute bgp 12703 metric transparent
!
router bgp 12703
address-family ipv4 vrf Customer-RIP-A1
redistribute rip
no auto-summary
```

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-8

The RIP configuration in this sample network (and described here) is very easy:

- The RIP routing process is configured. The RIP version is configured as the global RIP parameter.
- The RIP routing context is configured for every VRF where you want to run RIP as the PE-CE routing protocol. The directly connected networks (configured on interfaces in the VRF) over which you want to run RIP are specified to have standard RIP configuration.
- Redistribution from BGP into RIP with metric propagation is configured.
- The BGP routing context is configured for every VRF. Redistribution of RIP routes into BGP has to be configured for every VRF for which you have configured the RIP routing context.

Configuring EIGRP PE-CE Routing

This topic describes how to configure an EIGRP PE-CE routing session.

Configuring EIGRP PE-CE Routing

Cisco.com

- Provides EIGRP with the capability to redistribute routes through a VPN cloud
- Configuration of only the PE routers required
- No upgrade or configuration changes to customer equipment

Note: Because of current limitations with route redistribution, backdoor links are not supported.

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-9

MPLS VPN support for EIGRP between PE and CE provides EIGRP with the capability to redistribute routes through a BGP VPN cloud. This feature is configured only on PE routers, requiring no upgrade or configuration changes to customer equipment. This feature also introduces EIGRP support for MPLS and BGP extended community attributes.

Note Because of current limitations with route redistribution, backdoor links are not supported. If a backdoor link is implemented, it may become active and override selection of the VPN links.

Configuring EIGRP PE-CE Routing: EIGRP Metric Propagation

Cisco.com

Router (config)#

```
router eigrp process-id
  address-family ipv4 vrf vrf-name
    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 received without the configured metric are not to be advertised to the CE router.
 - The metric can be configured in the redistribute statement using the redistribute command or configured with the default-metric command.


© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-10

The IGP metric is always copied into the MED attribute of the BGP route when an IGP route is redistributed into BGP. Within standard BGP implementation, the MED attribute is used only as a route selection criterion. 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.

Configuring EIGRP PE-CE Routing: Example

Cisco.com



```
router eigrp 1
 address-family ipv4 vrf Customer-EIGRP-A1
  autonomous-system 101
  network 172.16.0.0 255.255.0.0
  redistribute BGP 12703 metric 10000 100 255 1 1500
  no auto-summary
!
router bgp 12703
 address-family ipv4 vrf Customer-EIGRP-A1
  redistribute eigrp 101
```

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-11

The EIGRP configuration in this sample network (and described here) is very easy:

- The EIGRP routing process is configured. The EIGRP process is configured as the global EIGRP parameter. Notice that this PE-CE configuration varies from traditional EIGRP configuration by deferring the definition of the AS number in the routing context.
- The EIGRP routing context is configured for every VRF where you want to run EIGRP as the PE-CE routing protocol. The directly connected networks (configured on interfaces in the VRF) over which you want to run EIGRP are specified to have standard EIGRP configuration.
- Redistribution from BGP into EIGRP with metric propagation is configured.
- The BGP routing context is configured for every VRF. Redistribution of EIGRP routes into BGP has to be configured for every VRF for which you have configured the EIGRP routing context.

Note Use of the **no auto-summary** command is recommended to prevent undesirable results in MPLS VPN. Use of the default **auto-summary** command can result in the same summary being received from multiple other sites based on network class and, therefore, you would not be able to determine which site to use for a more specific route.

Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- **The per-VRF routing protocols can be configured in two ways: as individual address families belonging to the same routing process or as separate routing processes.**
- **Use the `address-family ipv4 vrf vrf-name` command to select the VRF routing context.**
- **Use the `ip route vrf` command to establish static routes.**
- **Use the `address-family ipv4 vrf vrf-name` command to start the configuration of individual routing context.**
- **Use the `redistribute` command to configure the metric that is copied into the MED attribute of the BGP route.**

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-12

Monitoring MPLS VPN Operations

Overview

This lesson presents the commands, syntax, and descriptions for monitoring VRF routing, MP-BGP sessions, and VPN status.

It is important to understand the network that you configure and ensure that it is operating optimally. This lesson will explain how to monitor an MPLS VPN network to ensure that the network is functioning smoothly.

Objectives

Upon completing this lesson, you will be able to describe how to monitor MPLS VPN operations. This ability includes being able to meet these objectives:

- Monitor VRF information
- Monitor VRF routing
- Monitor MP-BGP sessions
- Monitor an MP-BGP VPNv4 table
- Monitor per-VRF CEF and LFIB structures
- Monitor labels associated with VPNv4 routes
- Identify the command syntax that is used with other MPLS VPN monitoring commands

Monitoring VRFs

This topic describes how to monitor VRF information.

Monitoring VRFs

Cisco.com

Router#
`show ip vrf`

- Displays the list of all VRFs configured in the router.

Router#
`show ip vrf detail`

- Displays detailed VRF configuration.

Router#
`show ip vrf interfaces`

- Displays interfaces associated with VRFs.

© 2004 Cisco Systems, Inc. All rights reserved.MPLS v2.1—5-3

show ip vrf

To display the set of defined VRFs and associated interfaces, use the **show ip vrf** command in EXEC mode: **show ip vrf** [{**brief** | **detail** | **interfaces**}] [*vrf-name*] [*output-modifiers*].

This table describes the parameters for the **show ip vrf** command.

Syntax Description

Parameter	Description
brief	(Optional) Displays concise information on the VRF (or VRFs) and associated interfaces.
detail	(Optional) Displays detailed information on the VRF (or VRFs) and associated interfaces.
interfaces	(Optional) Displays detailed information about all interfaces bound to a particular VRF or to any VRF.
<i>vrf-name</i>	(Optional) Name assigned to a VRF.
<i>output-modifiers</i>	(Optional) For a list of associated keywords and arguments, use context-sensitive help.

Defaults

When no optional parameters are specified, the command shows concise information about all configured VRFs.

Monitoring VRFs: show ip vrf

Cisco.com

```
Router#show ip vrf
  Name          Default RD      Interfaces
SiteA2         103:30         Serial1/0.20
SiteB          103:11         Serial1/0.100
SiteX          103:20         Ethernet0/0
Router#
```

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-4

The **show ip vrf** command displays concise information about the VRF (or VRFs) and associated interfaces. This table describes the fields displayed by this command.

Field Description

Fields	Description
Name	Specifies the VRF name.
Default RD	Specifies the default RD.
Interfaces	Specifies the network interfaces.

Monitoring VRFs: show ip vrf detail

Cisco.com

```
Router#show ip vrf detail
VRF SiteA2; default RD 103:30
  Interfaces:
    Serial1/0.20
    Connected addresses are not in global routing table
    No Export VPN route-target communities
    Import VPN route-target communities
      RT:103:10
    No import route-map
    Export route-map: A2
VRF SiteB; default RD 103:11
  Interfaces:
    Serial1/0.100
    Connected addresses are not in global routing table
    Export VPN route-target communities
      RT:103:11
    Import VPN route-target communities
      RT:103:11          RT:103:20
    No import route-map
    No export route-map
```

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-5

To display detailed information on the VRFs and associated interfaces, use the **show ip vrf detail** command. This table describes the additional fields shown by this command.

Additional Field Description

Field	Description
Interfaces	Specifies the network interfaces.
Export	Specifies VPN RT export communities.
Import	Specifies VPN RT import communities.

Monitoring VRFs: show ip vrf interfaces

Cisco.com

```
Router#show ip vrf interfaces
Interface      IP-Address      VRF      Protocol
Serial1/0.20   150.1.31.37     SiteA2   up
Serial1/0.100  150.1.32.33     SiteB    up
Ethernet0/0    192.168.22.3    SiteX    up
```

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-6

To display the interfaces bound to a particular VRF (or interfaces bound to any VRF), use the **show ip vrf interfaces** command, which displays the fields described in this table.

show ip vrf interfaces Field Description

Field	Description
Interface	Specifies the network interfaces for a VRF.
IP-Address	Specifies the IP address of a VRF interface.
VRF	Specifies the VRF name.
Protocol	Displays the state of the protocol (up or down) for each VRF interface.

Monitoring VRF Routing

This topic describes how to monitor VRF routing.

Monitoring VRF Routing

Cisco.com

Router#
`show ip protocols vrf name`

- Displays the routing protocols configured in a VRF.

Router#
`show ip route vrf name`

- Displays the VRF routing table.

Router#
`show ip bgp vpnv4 vrf name`

- Displays per-VRF BGP parameters.

© 2004 Cisco Systems, Inc. All rights reserved.MPLS v2.1—5-7

The following three commands can be used to monitor VRF routing:

- The **show ip protocols vrf** command displays the summary information about routing protocols running in a VRF.
- The **show ip route vrf** command displays the VRF routing table.
- The **show ip bgp vpnv4 vrf** command displays the VRF BGP table.

show ip protocols vrf

To display the routing protocol information associated with a VRF, use the **show ip protocols vrf** command in EXEC mode: **show ip protocols vrf vrf-name**.

This table describes the parameters for the **show ip protocols vrf** command.

Syntax Description

Parameter	Description
<i>vrf-name</i>	Specifies name assigned to the VRF.

show ip route vrf

To display the IP routing table associated with a VRF, use the **show ip route vrf** command in EXEC mode: **show ip route vrf** *vrf-name* [**connected**] [*protocol* [*as-number*] [*tag*] [*output-modifiers*]] [**list number** [*output-modifiers*]] [**profile**] [**static** [*output-modifiers*]] [**summary** [*output-modifiers*]] [**supernets-only** [*output-modifiers*]] [**traffic-engineering** [*output-modifiers*]].

This table describes the parameters for the **show ip route vrf** command.

Syntax Description

Parameter	Description
<i>vrf-name</i>	Specifies name assigned to the VRF.
connected	(Optional) Displays all connected routes in a VRF.
<i>protocol</i>	(Optional) To specify a routing protocol, use one of the following keywords: bgp , egp , eigrp , hello , igrp , isis , ospf , or rip .
<i>as-number</i>	(Optional) Specifies AS number.
<i>tag</i>	(Optional) Specifies Cisco IOS software routing area label.
<i>output-modifiers</i>	(Optional) For a list of associated keywords and arguments, use context-sensitive help.
list number	(Optional) Specifies the IP access list to display.
profile	(Optional) Displays the IP routing table profile.
static	(Optional) Displays static routes.
summary	(Optional) Displays a summary of routes.
supernets-only	(Optional) Displays supernet entries only.
traffic-engineering	(Optional) Displays only traffic-engineered routes.

show ip bgp vpnv4

To display VPN address information from the BGP table, use the **show ip bgp vpnv4** command in EXEC mode: **show ip bgp vpnv4** {**all** | **rd** *route-distinguisher* | **vrf** *vrf-name*} [*ip-prefix/length*] [**longer-prefixes**] [*output-modifiers*] [*network-address* [*mask*]] [**longer-prefixes**] [*output-modifiers*] [**cidr-only**] [**community**] [**community-list**] [**dampened-paths**] [**filter-list**] [**flap-statistics**] [**inconsistent-as**] [**neighbors**] [**paths** [*line*]] [**peer-group**] [**quote-regexp**] [**regexp**] [**summary**] [**labels**].

This table describes the parameters for the **show ip bgp vpnv4** command.

Syntax Description

Parameter	Description
all	Displays the complete VPNv4 database.
rd <i>route-distinguisher</i>	Displays Network Layer Reachability Information (NLRI) prefixes that have a matching RD.
vrf <i>vrf-name</i>	Displays NLRI prefixes associated with the named VRF.
<i>ip-prefix/length</i>	(Optional) IP prefix address (in dotted decimal notation) and length of mask (0 to 32).
longer-prefixes	(Optional) Displays the entry, if any, that exactly matches the specified prefix parameter, and all entries that match the prefix in a “longest-match” sense—that is, prefixes for which the specified prefix is an initial substring.
<i>output-modifiers</i>	(Optional) For a list of associated keywords and arguments, use context-sensitive help.
<i>network-address</i>	(Optional) IP address of a network in the BGP routing table.
<i>mask</i>	(Optional) Mask of the network address, in dotted decimal notation.
cidr-only	(Optional) Displays only routes that have non-natural net masks.
community	(Optional) Displays routes matching this community.
community-list	(Optional) Displays routes matching this community list.
dampened-paths	(Optional) Displays paths suppressed on account of dampening (BGP route from peer is up and down).
filter-list	(Optional) Displays routes conforming to the filter list.
flap-statistics	(Optional) Displays flap statistics of routes.
inconsistent-as	(Optional) Displays only routes that have inconsistent autonomous systems of origin.
neighbors	(Optional) Displays details about TCP and BGP neighbor connections.
paths	(Optional) Displays path information.
<i>line</i>	(Optional) A regular expression to match the BGP AS paths.
peer-group	(Optional) Displays information about peer groups.
quote-regexp	(Optional) Displays routes matching the AS path “regular expression.”

Parameter	Description
regex	(Optional) Displays routes matching the AS path "regular expression."
summary	(Optional) Displays BGP neighbor status.
tags	(Optional) Displays incoming and outgoing BGP labels for each NLRI.

Monitoring VRF Routing: show ip protocols vrf

Cisco.com

```

Router#show ip protocol vrf SiteX
Routing Protocol is "rip"
  Sending updates every 30 seconds, next due in 10 seconds
  Invalid after 180 seconds, hold down 180, flushed after 240
  Outgoing update filter list for all interfaces is
  Incoming update filter list for all interfaces is
  Redistributing: rip, bgp 3
  Default version control: send version 2, receive version 2
    Interface          Send Recv Triggered RIP Key-chain
  Ethernet0/0         2      2
  Routing for Networks:
    192.168.22.0
  Routing Information Sources:
    Gateway           Distance      Last Update
  Distance: (default is 120)

```

© 2004 Cisco Systems, Inc. All rights reserved.
MPLS v2.1—5-8

The **show ip protocols vrf** command displays summary information about all routing protocol instances active in the specified VRF. The fields displayed by this command are shown in this table.

Field Description

Field	Description
Gateway	Displays the IP address of the router identifier for all routers in the network.
Distance	Displays the metric used to access the destination route.
Last Update	Displays the last time that the routing table was updated from the source.

Monitoring VRF Routing: show ip route vrf

Cisco.com

```
Router#show ip route vrf SiteA2
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
        D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
        N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
        E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
        i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
        * - candidate default, U - per-user static route, o - ODR
        P - periodic downloaded static route

Gateway of last resort is not set

O    203.1.20.0/24 [110/782] via 150.1.31.38, 02:52:13, Serial1/0.20
     203.1.2.0/32 is subnetted, 1 subnets
O    203.1.2.1 [110/782] via 150.1.31.38, 02:52:13, Serial1/0.20
     203.1.1.0/32 is subnetted, 1 subnets
B    203.1.1.1 [200/1] via 192.168.3.103, 01:14:32
B    203.1.135.0/24 [200/782] via 192.168.3.101, 02:05:38
B    203.1.134.0/24 [200/1] via 192.168.3.101, 02:05:38
B    203.1.10.0/24 [200/1] via 192.168.3.103, 01:14:32

... rest deleted ...
```

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-9

The **show ip route vrf** command displays the contents of the VRF IP routing table in the same format used by the **show ip route** command.

Monitoring VRF Routing: show ip bgp vpnv4 vrf neighbors

Cisco.com

```
Router#show ip bgp vpnv4 vrf SiteB neighbors
BGP neighbor is 150.1.32.34, vrf SiteB, remote AS 65032, external link
BGP version 4, remote router ID 203.2.10.1
BGP state = Established, up for 02:01:41
Last read 00:00:56, hold time is 180, keepalive interval is 60 seconds
Neighbor capabilities:
  Route refresh: advertised and received
  Address family IPv4 Unicast: advertised and received
Received 549 messages, 0 notifications, 0 in queue
Sent 646 messages, 0 notifications, 0 in queue
Route refresh request: received 0, sent 0
Minimum time between advertisement runs is 30 seconds

For address family: VPNv4 Unicast
Translates address family IPv4 Unicast for VRF SiteB
BGP table version 416, neighbor version 416
Index 4, Offset 0, Mask 0x10
Community attribute sent to this neighbor
2 accepted prefixes consume 120 bytes
Prefix advertised 107, suppressed 0, withdrawn 63

... rest deleted ...
```

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-10

show ip bgp vpnv4 vrf neighbors

To display BGP neighbors configured in a VRF, use the **show ip bgp vpnv4 vrf neighbors** command in privileged EXEC mode: **show ip bgp vpnv4 {all | vrf vrf-name} neighbors**.

This table describes the parameters for the **show ip bgp vpnv4 vrf neighbors** command.

Syntax Description

Parameter	Description
vpn4	Specifies VPNv4 information.
all	Displays the complete VPNv4 database.
vrf vrf-name	Displays neighbors associated with the named VRF.
neighbors	Displays details about TCP and BGP neighbor connections.

Defaults

This command has no default values.

Usage Guidelines

Use this command to display detailed information about BGP neighbors associated with the MPLS VPN architecture.

Monitoring MP-BGP Sessions

This topic describes how to monitor MP-BGP sessions.

Monitoring MP-BGP Sessions

Cisco.com

```
Router#  
show ip bgp neighbors
```

- This command displays global BGP neighbors and the protocols negotiated with these neighbors.

© 2004 Cisco Systems, Inc. All rights reserved. MPLS v2.1—5-11

The **show ip bgp neighbors** command is described in detail in Cisco IOS documentation. This command is used to monitor BGP sessions with other PE routers and the address families negotiated with these neighbors.

Monitoring MP-BGP Sessions: show ip bgp neighbors

Cisco.com

```
Router#show ip bgp neighbor 192.168.3.101  
BGP neighbor is 192.168.3.101, remote AS 3, internal link  
BGP version 4, remote router ID 192.168.3.101  
BGP state = Established, up for 02:15:33  
Last read 00:00:33, hold time is 180, keepalive interval is 60 seconds  
Neighbor capabilities:  
  Route refresh: advertised and received  
  Address family IPv4 Unicast: advertised and received  
  Address family VPNv4 Unicast: advertised and received  
Received 1417 messages, 0 notifications, 0 in queue  
Sent 1729 messages, 2 notifications, 0 in queue  
Route refresh request: received 9, sent 29  
Minimum time between advertisement runs is 5 seconds  
  
For address family: IPv4 Unicast  
BGP table version 188, neighbor version 188  
Index 2, Offset 0, Mask 0x4  
1 accepted prefixes consume 36 bytes  
Prefix advertised 322, suppressed 0, withdrawn 230  
  
... Continued
```

© 2004 Cisco Systems, Inc. All rights reserved. MPLS v2.1—5-12

Monitoring MP-BGP Sessions: show ip bgp neighbors (Cont.)

Cisco.com

```
Router#show ip bgp neighbor 192.168.3.101

... Continued

For address family: VPNv4 Unicast
BGP table version 416, neighbor version 416
Index 2, Offset 0, Mask 0x4
NEXT_HOP is always this router
Community attribute sent to this neighbor
6 accepted prefixes consume 360 bytes
Prefix advertised 431, suppressed 0, withdrawn 113

Connections established 7; dropped 6
Last reset 02:18:33, due to Peer closed the session

... Rest deleted
```

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-13

show ip bgp neighbors

To display information about the TCP and BGP connections to neighbors, use the **show ip bgp neighbors** command in EXEC mode: **show ip bgp neighbors** [*neighbor-address*] [**received-routes** | **routes** | **advertised-routes** | {*paths regexp*} | **dampened-routes**].

This table describes the parameters for the **show ip bgp neighbors** command.

Syntax Description

Parameter	Description
<i>neighbor-address</i>	(Optional) Address of the neighbor whose routes you have learned from. If you omit this argument, all neighbors will be displayed.
received-routes	(Optional) Displays all received routes (both accepted and rejected) from the specified neighbor.
routes	(Optional) Displays all routes that are received and accepted. This parameter is a subset of the output from the received-routes keyword.
advertised-routes	(Optional) Displays all the routes that the router has advertised to the neighbor.
paths <i>regexp</i>	(Optional) Matches the paths received.
dampened-routes	(Optional) Displays the dampened routes to the neighbor at the IP address specified.

Example: Sample Output from show ip bgp neighbors Command

This table describes the fields shown in the sample output.

Field Descriptions

Field	Description
BGP neighbor	IP address of the BGP neighbor and its AS number. If the neighbor is in the same AS as the router, the link between them is internal; otherwise, the link is considered external.
remote AS	AS of the neighbor.
external link internal link	Indicates that this peer is either an EBGP peer or an IBGP peer.
BGP version	BGP version being used to communicate with the remote router. The router ID (an IP address) of the neighbor is also specified.
remote router ID	IP address of the neighbor.
BGP state	Internal state of this BGP connection.
up for	Amount of time, in seconds, that the underlying TCP connection has been in existence.
Last read	Time that BGP last read a message from this neighbor.
hold time	Maximum amount of time that can elapse between messages from the peer.
keepalive interval	Time period, in seconds, between sending keepalive packets, which helps ensure that the TCP connection is up.
Neighbor capabilities	BGP capabilities advertised and received from this neighbor. Note: A state of "advertised and received" indicates an active neighbor relationship.

The following is sample output from the **show ip bgp neighbors** command:

```
Router# sh ip bgp nei 192.168.100.129
BGP neighbor is 192.168.100.129, remote AS 65001, internal link
BGP version 4, remote router ID 192.168.100.129
BGP state = Established, up for 5d01h
Last read 00:00:56, hold time is 180, keepalive interval is 60 seconds
Neighbor capabilities:
  Route refresh: advertised and received(old & new)
  Address family IPv4 Unicast: advertised and received
  Address family VPNv4 Unicast: advertised and received

For address family: IPv4 Unicast
BGP table version 31, neighbor version 31
Index 1, Offset 0, Mask 0x2
Prefix activity:
  Sent      Rcvd
Prefixes Current:      0      30 (Consumes 1440 bytes)
Prefixes Total:        0      30
```

```

Implicit Withdraw:          0          0
Explicit Withdraw:        0          0
Used as bestpath:         n/a        30
Used as multipath:        n/a        0
                          Outbound   Inbound
Local Policy Denied Prefixes: -----
  Bestpath from this peer:          30          n/a
  Total:                            30          0
Number of NLRI in the update sent: max 0, min 0

```

```

For address family: VPNv4 Unicast
BGP table version 30, neighbor version 30
Index 4, Offset 0, Mask 0x10
NEXT_HOP is always this router
Community attribute sent to this neighbor

```

```

                          Sent      Rcvd
Prefix activity:         ----      ----
  Prefixes Current:      9          1 (Consumes 256 bytes)
  Prefixes Total:        18         1
  Implicit Withdraw:     9          0
  Explicit Withdraw:     0          0
  Used as bestpath:      n/a        4
  Used as multipath:     n/a        0

```

```

                          Outbound   Inbound
Local Policy Denied Prefixes: -----
  VPN Imported prefix:          3          n/a
  Bestpath from this peer:      1          n/a
  Total:                        4          0
Number of NLRI in the update sent: max 4, min 0

```

(output omitted)

This table describes the fields shown in the sample output.

Field Descriptions

Field	Description
Address family IPv4 Unicast:	IPv4 unicast-specific properties of this neighbor.
Address family VPNv4	VPNv4-specific properties of this neighbor.

Note For detailed information, please consult the Cisco IOS reference manual.

Monitoring an MP-BGP VPNv4 Table

This topic describes how to monitor an MP-BGP VPNv4 table.

Monitoring an MP-BGP VPNv4 Table

Cisco.com

Router#
`show ip bgp vpnv4 all`

- Displays whole VPNv4 table.

Router#
`show ip bgp vpnv4 vrf vrf-name`

- Displays only BGP parameters (routes or neighbors) associated with specified VRF.
- Any BGP show command can be used with these parameters.

Router#
`show ip bgp vpnv4 rd route-distinguisher`

- Displays only BGP parameters (routes or neighbors) associated with specified RD.

© 2004 Cisco Systems, Inc. All rights reserved. MPLS v2.1—5-14

The **show ip bgp vpnv4** command displays IPv4 BGP information and VPNv4 BGP information. To display VPNv4 BGP information, use one of these keywords:

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

Monitoring an MP-BGP VPNv4 Table: show ip bgp vpnv4 vrf-name

Cisco.com

```

Router#show ip bgp vpnv4 vrf SiteA2
BGP table version is 416, local router ID is 192.168.3.102
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
Route Distinguisher: 103:30 (default for vrf SiteA2)
*> 150.1.31.36/30   0.0.0.0           0         32768 ?
*>i150.1.31.128/30  192.168.3.101    0         100    0 ?
*>i150.1.31.132/30  192.168.3.101    0         100    0 ?
*>i203.1.1.1/32     192.168.3.103    1         100    0 65031 i
*> 203.1.2.1/32    150.1.31.38      782        32768 ?
*>i203.1.10.0      192.168.3.103    1         100    0 65031 i
*> 203.1.20.0      150.1.31.38      782        32768 ?
*>i203.1.127.3/32  192.168.3.101    1         100    0 ?
*>i203.1.127.4/32  192.168.3.101    782        100    0 ?
*>i203.1.134.0     192.168.3.101    1         100    0 ?
*>i203.1.135.0     192.168.3.101    782        100    0 ?
  
```

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-15

show ip bgp vpnv4 vrf

To display VPNv4 information from the BGP database associated with a VRF, use the **show ip bgp vpnv4 vrf** command in privileged EXEC mode: **show ip bgp vpnv4 vrf vrf-name** [*ip-prefix/length*] [**longer-prefixes**] [*output-modifiers*] [*network-address* [*mask*]] [**longer-prefixes**] [*output-modifiers*] [**cidr-only**] [**community**][**community-list**] [**dampened-paths**] [**filter-list**] [**flap-statistics**] [**inconsistent-as**] [**neighbors**] [**paths** [*line*]] [**peer-group**] [**quote-regexp**] [**regexp**] [**summary**] [**tags**].

This table describes the syntax for the **show ip bgp vpnv4 vrf** command.

Syntax Description

Parameter	Description
vrf <i>vrf-name</i>	Displays NLRI prefixes associated with the named VRF.

Defaults

This command has no default values.

Usage Guidelines

Use this command to display VPNv4 information that is associated with a VRF from the BGP database. A similar command—**show ip bgp vpnv4 all**—displays all available VPNv4 information. The **show ip bgp vpnv4 summary** command displays BGP neighbor status.

Monitoring an MP-BGP VPNv4 Table: show ip bgp vpnv4 rd *route-distinguisher*

Cisco.com

```
Router#show ip bgp vpnv4 rd 103:30 203.1.127.3
BGP routing table entry for 103:30:203.1.127.3/32, version 164
Paths: (1 available, best #1, table SiteA2)
  Not advertised to any peer
  Local, imported path from 103:10:203.1.127.3/32
    192.168.3.101 (metric 10) from 192.168.3.101 (192.168.3.101)
      Origin incomplete, metric 1, localpref 100, valid,
        internal, best
    Extended Community: RT:103:10
```

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-16

show ip bgp vpnv4 rd *route-distinguisher*

To display all VPNv4 routes that contain a specified RD, use the **show ip bgp vpnv4 rd** command in privileged EXEC mode: **show ip bgp vpnv4 rd *route-distinguisher*** [*ip-prefix/length*] [**longer-prefixes**] [*output-modifiers*] [*network-address*] [*mask*] [**longer-prefixes**] [*output-modifiers*] [**cidr-only**] [**community**][**community-list**] [**dampened-paths**] [**filter-list**] [**flap-statistics**] [**inconsistent-as**] [**paths**] [*line*] [**quote-regexp**] [**regexp**] [**summary**] [**tags**].

This table describes the syntax for the **show ip bgp vpnv4 rd *route-distinguisher*** command.

Syntax Description

Parameter	Description
rd <i>route-distinguisher</i>	Displays NLRI prefixes that have a matching RD.

Defaults

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

Usage Guidelines

An RD creates routing and forwarding tables and specifies the default RD for a VPN. The RD is added to the beginning of the customer IPv4 prefixes to change them into globally unique VPN IPv4 prefixes.

Either an RD is an AS number relative RD, in which case it is composed of an AS number and an arbitrary number, or it is an IP-address-relative RD, in which case it is composed of an IP address and an arbitrary number.

You can enter an RD in either of these formats:

- *16-bit AS number:* your 32-bit number
 - For example, 101:3.
- *32-bit IP address:* your 16-bit number
 - For example, 192.168.122.15:1.

Example: Configuring a Default RD for Two VRFs

The following example shows how to configure a default RD for two VRFs. The example illustrates the use of both AS-relative and IP-address-relative RDs:

```
Router(config)# ip vrf vrf_blue
Router(config-vrf)# rd 100:3
Router (config-vrf)# exit
Router(config)# ip vrf vrf_red
Router(config-vrf)# rd 173.13.0.12:200
```

Monitoring Per-VRF CEF and LFIB Structures

This topic describes how to monitor per-VRF CEF and label forwarding information base (LFIB) structures.

Monitoring per-VRF CEF and LFIB Structures

Cisco.com

Router#
`show ip cef vrf vrf-name`

- Displays per-VRF CEF table.

Router#
`show ip cef vrf vrf-name ip-prefix detail`

- Displays details of an individual CEF entry, including label stack.

Router#
`show mpls forwarding vrf vrf-name`

- Displays labels allocated by an MPLS VPN for routes in the specified VRF.

© 2004 Cisco Systems, Inc. All rights reserved.MPLS v2.1—5-17

The following three commands can be used to display per-VRF FIB and LFIB structures:

- The **show ip cef vrf** command displays the VRF FIB.
- The **show ip cef vrf detail** command displays detailed information about a single entry in the VRF FIB.
- The **show mpls forwarding vrf** command displays all labels allocated to VPN routes in the specified VRF.

Monitoring per-VRF CEF and LFIB Structures (Cont.)

Cisco.com

```
Router#show ip cef vrf SiteA2 203.1.1.1 255.255.255.255 detail
203.1.1.1/32, version 57, cached adjacency to Serial1/0.2
0 packets, 0 bytes
tag information set
  local tag: VPN-route-head
  fast tag rewrite with S1/0.2, point2point, tags imposed: {26 39}
via 192.168.3.103, 0 dependencies, recursive
next hop 192.168.3.10, Serial1/0.2 via 192.168.3.103/32
valid cached adjacency
tag rewrite with S1/0.2, point2point, tags imposed: {26 39}
```

The `show ip cef` command can also display the label stack associated with the MP-IBGP route.

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-18

show ip cef vrf

To display the CEF forwarding table associated with a VRF, use the `show ip cef vrf` command in privileged EXEC mode: `show ip cef vrf vrf-name [ip-prefix [mask [longer-prefixes]] [detail] [output-modifiers]] [interface interface-number] [adjacency [interface interface-number] [detail] [discard] [drop] [glean] [null] [punt] [output-modifiers]] [detail [output-modifiers]] [non-recursive [detail] [output-modifiers]] [summary [output-modifiers]] [traffic [prefix-length] [output-modifiers]] [unresolved [detail] [output-modifiers]]`.

This table describes the syntax for the **show ip cef vrf** command.

Syntax Description

Parameter	Description
<i>vrf-name</i>	Name assigned to the VRF.
<i>ip-prefix</i>	(Optional) IP prefix of entries to show, in dotted decimal notation (A.B.C.D).
<i>mask</i>	(Optional) Mask of the IP prefix, in dotted decimal notation.
longer-prefixes	(Optional) Displays table entries for all of the more specific routes.
detail	(Optional) Displays detailed information for each CEF table entry.
<i>output-modifiers</i>	(Optional) For a list of associated keywords and arguments, use context-sensitive help.
<i>interface</i>	(Optional) Type of network interface to use: ATM, Ethernet, loopback, POS, or null.
<i>interface-number</i>	Number identifying the network interface to use.
adjacency	(Optional) Displays all prefixes resolving through adjacency.
discard	Discards adjacency.
drop	Drops adjacency.
glean	Gleans adjacency.
null	Null adjacency.
punt	Punts adjacency.
non-recursive	(Optional) Displays only nonrecursive routes.
summary	(Optional) Displays a CEF table summary.
traffic	(Optional) Displays traffic statistics.
prefix-length	(Optional) Displays traffic statistics by prefix size.
unresolved	(Optional) Displays only unresolved routes.

Defaults

This command has no default values.

Usage Guidelines

Used with the *vrf-name* argument, the **show ip cef vrf** command shows a shortened display of the CEF table.

Used with the **detail** keyword, the **show ip cef vrf** command shows detailed information for all CEF table entries.

Monitoring per-VRF CEF and LFIB Structures (Cont.)

Cisco.com

```
Router#show mpls forwarding vrf SiteA2
Local  Outgoing  Prefix          Bytes tag  Outgoing  Next Hop
tag    tag or VC    or Tunnel Id    switched   interface
26     Aggregate   150.1.31.36/30 [V] 0
37     Untagged    203.1.2.1/32 [V] 0          Se1/0.20  point2point
38     Untagged    203.1.20.0/24 [V] 0          Se1/0.20  point2point

Router#show mpls forwarding vrf SiteA2 tags 37 detail
Local  Outgoing  Prefix          Bytes tag  Outgoing  Next Hop
tag    tag or VC    or Tunnel Id    switched   interface
37     Untagged    203.1.2.1/32 [V] 0          Se1/0.20  point2point
      MAC/Encaps=0/0, MTU=1504, Tag Stack{}
      VPN route: SiteA2
      Per-packet load-sharing
```

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-19

show mpls forwarding vrf

To display label-forwarding information for advertised VRF routes, use the **show mpls forwarding vrf** command in EXEC mode. To disable the display of label-forwarding information, use the **no** form of this command.

- **show mpls forwarding vrf** *vrf-name* [*ip-prefix/length* [*mask*]] [**detail**] [*output-modifiers*]
- **no show mpls forwarding vrf** *vrf-name* [*ip-prefix/length* [*mask*]] [*detail*] [*output-modifiers*]

This table describes the parameters for the **show mpls forwarding vrf** command.

Syntax Description

Parameter	Description
<i>vrf-name</i>	Displays NLRI prefixes associated with the named VRF.
<i>ip-prefix/length</i>	(Optional) IP prefix address (in dotted decimal notation) and length of mask (0 to 32).
<i>mask</i>	(Optional) Destination network mask in dotted decimal notation.
detail	(Optional) Displays detailed information on the VRF routes.
<i>output-modifiers</i>	(Optional) For a list of associated keywords and arguments, use context-sensitive help.

Defaults

This command has no default behavior or values.

Usage Guidelines

Use this command to display label-forwarding entries associated with a particular VRF or IP prefix.

Monitoring Labels Associated with VPNv4 Routes

This topic describes how to monitor labels associated with VPNv4 routes.

Monitoring Labels Associated with VPNv4 Routes

Cisco.com

Router#

```
show ip bgp vpnv4 [ all | rd value | vrf-name ] tags
```

- Displays labels associated with VPNv4 routes.

```
Router#show ip bgp vpnv4 all tags
```

Network	Next Hop	In tag/Out tag
Route Distinguisher: 100:1 (vrf1)		
2.0.0.0	10.20.0.60	34/notag
10.0.0.0	10.20.0.60	35/notag
12.0.0.0	10.20.0.60	26/notag
	10.20.0.60	26/notag
13.0.0.0	10.15.0.15	notag/26

© 2004 Cisco Systems, Inc. All rights reserved.MPLS v2.1—5-20

You can use the **show ip bgp vpnv4 tags** command to display tags assigned to local or remote VRF routes by the local or remote PE router. This command displays tags associated with all VPNv4 routes when you use the **all** keyword. This command can also display tags associated with a specified RD or VRF.

This table describes the fields displayed by this command.

Field Descriptions

Field	Description
Network	Displays the network address from the BGP table.
Next Hop	Specifies the BGP next-hop address.
In tag	Displays the label (if any) assigned by this router.
Out tag	Displays the label assigned by the BGP next-hop router.

Identifying Other MPLS VPN Monitoring Commands

This topic identifies the command syntax that is used with other MPLS VPN monitoring commands.

Other MPLS VPN Monitoring Commands

Cisco.com

Router#
`telnet host /vrf vrf-name`

- Performs PE-CE Telnet through specified VRF.

Router#
`ping vrf vrf-name ...`

- Performs ping based on VRF routing table.

Router#
`trace vrf vrf-name ...`

- Performs VRF-based traceroute.

© 2004 Cisco Systems, Inc. All rights reserved.MPLS v2.1—5-21

The following three additional Cisco IOS software monitoring commands are VRF-aware:

- The **telnet** command can be used to connect to a CE router from a PE router using the **/vrf** option.
- The **ping vrf** command can be used to ping a destination host reachable through a VRF.
- The **trace vrf** command can be used to trace a path toward a destination reachable through a VRF.

Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- **Use the following commands to monitor VRF information:**
 - show ip vrf
 - show ip vrf detail
 - show ip vrf interfaces
- **Use the following commands to monitor VRF routing:**
 - show ip protocols vrf
 - show ip route vrf
 - show ip bgp vpnv4 vrf
- **Use the show ip bgp neighbors command to monitor MP-BGP sessions.**

© 2004 Cisco Systems, Inc. All rights reserved. MPLS v2.1—5-22

Summary (Cont.)

Cisco.com

- **Use the show ip bgp vpnv4 command to monitor an MP-BGP VPNv4 table.**
- **Use the show ip bgp vpnv4 all tags command to monitor MP-BGP VPNv4 labels.**
- **Use the following commands to monitor the per-VRF CEF and LFIB structures:**
 - show ip cef vrf
 - show ip cef vrf detail
 - show mpls forwarding vrf
- **Other commands to monitor MPLS VPN are as follows:**
 - telnet
 - ping vrf
 - trace vrf

© 2004 Cisco Systems, Inc. All rights reserved. MPLS v2.1—5-23

Configuring OSPF as the Routing Protocol Between PE and CE Routers

Overview

This lesson explains the PE-CE routing protocol configuration steps required when you are running OSPF between PE and CE routers, and the issues that may be encountered.

It is important to understand not only what you can configure between PE and CE routers when you are setting up MPLS VPNs, but also how to accomplish the configuration successfully. This lesson looks at the configuration parameters that you need to configure an MPLS VPN PE-CE routing exchange.

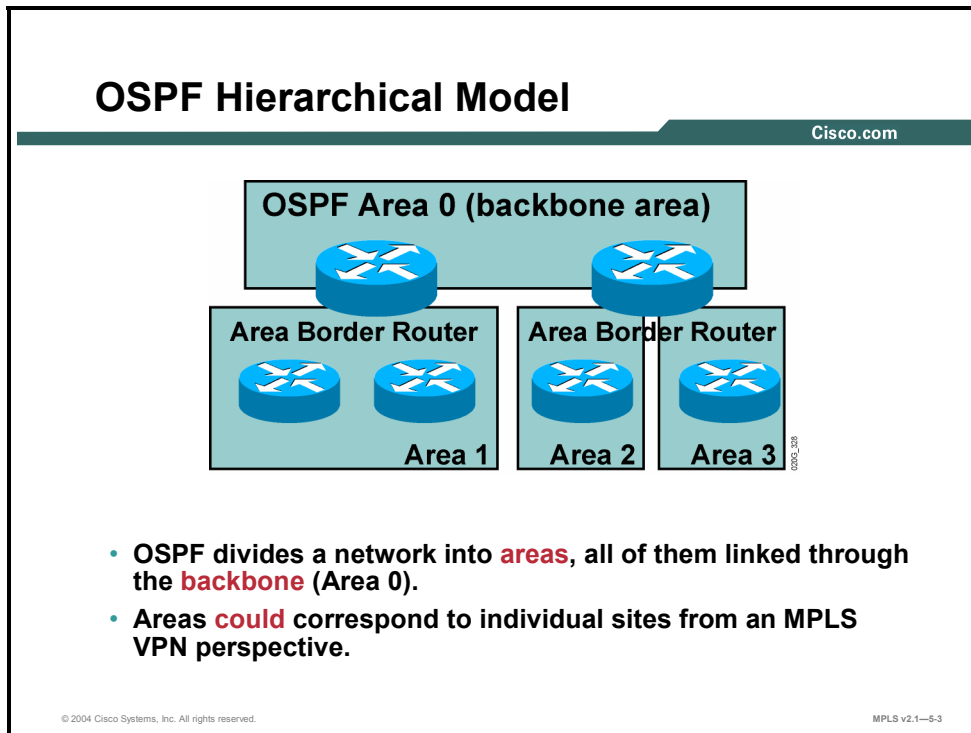
Objectives

Upon completing this lesson, you will be able to describe how to configure OSPF as the routing protocol between PE-CE routers. This ability includes being able to meet these objectives.

- Describe the features of the OSPF hierarchical model
- Describe the propagation of OSPF customer routes across the MPLS VPN backbone
- Describe how an MPLS VPN is implemented as an OSPF superbackbone
- Describe how to configure a PE-CE OSPF routing session
- Describe how the OSPF down bit is used to address the route loop issue
- Describe how packet forwarding is optimized across the MPLS VPN
- Describe how the OSPF tag field is used to address the root loop issue
- Describe the features of a sham link
- Describe how to configure a sham link

What Is the Enhanced OSPF Hierarchical Model?

This topic describes the features of the enhanced OSPF hierarchical model.



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


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

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

Note Please refer to the *Building Scalable Cisco Internetworks* (BSCI) course for background information on OSPF.

Propagating OSPF Customer Routes


This topic describes the propagation of OSPF customer routes across the MPLS VPN backbone.



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

OSPF in an MPLS VPN Routing Model: OSPF-BGP Redistribution Issue

Cisco.com



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

OSPF in an MPLS VPN Routing Model: Classic OSPF-BGP Redistribution

Cisco.com

- OSPF route type is not preserved when OSPF route is redistributed into BGP.
- All OSPF routes from a site are inserted as external (type 5 LSA) routes into other sites.
- **Result:** OSPF route summarization and stub areas are hard to implement.
- **Conclusion:** MPLS VPN must extend the classic OSPF-BGP routing model.

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-6

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

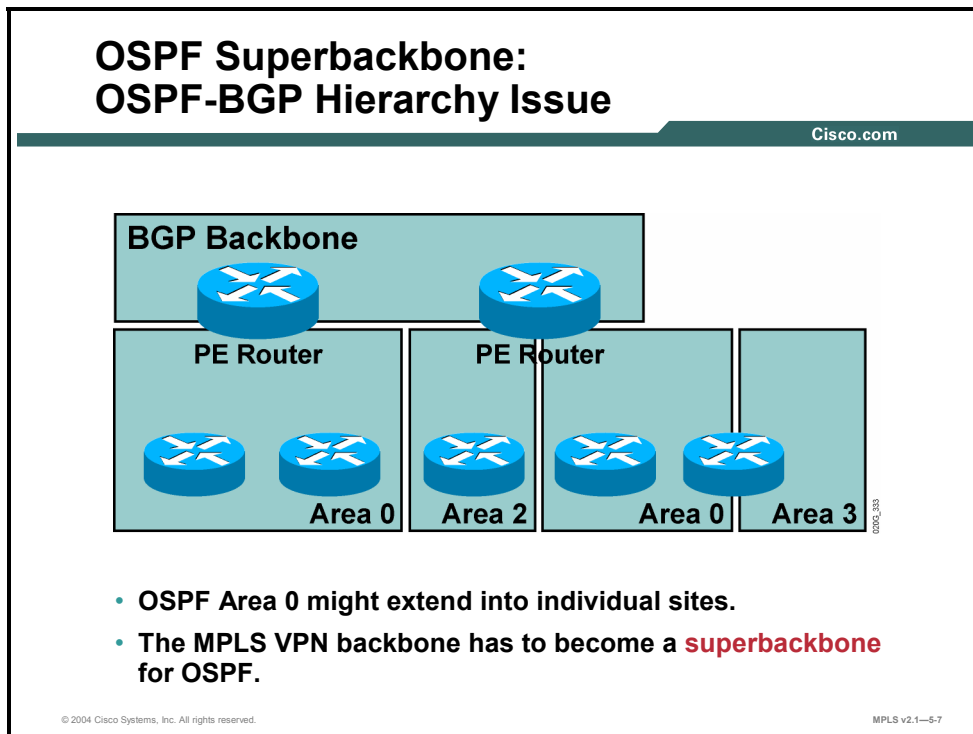
The following identifies some of the caveats associated with external OSPF routes:

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

Because of all these caveats, migrating an OSPF customer toward an MPLS VPN service might 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.

Implementing MPLS VPNs as an OSPF Superbackbone

This topic describes how an MPLS VPN is implemented as an OSPF superbackbone.



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

OSPF in MPLS VPNs: Goals

Cisco.com

- **OSPF between sites shall not use normal OSPF-BGP redistribution.**
- **OSPF continuity must be provided across 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.**

© 2004 Cisco Systems, Inc. All rights reserved.


MPLS v2.1—5-8

Here are the goals that have to be met by the OSPF superbackbone:

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

OSPF Superbackbone: Route Propagation Example

Cisco.com



© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-9

The MPLS VPN superbackbone appears as another layer of hierarchy in the OSPF architecture. The PE routers that connect regular OSPF areas to the superbackbone therefore appear as OSPF Area Border Routers (ABRs) in the OSPF areas to which they are attached. In Cisco IOS implementation, 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 following describes how the continuity of OSPF routing is preserved:

- The OSPF intra-area route—described in the OSPF router link-state advertisement (LSA) or network LSA—is inserted into the OSPF superbackbone by redistributing the OSPF route into MP-BGP. Route summarization can be performed on the redistribution boundary by the PE router.
- The MP-BGP route is propagated to other PE routers and inserted as an OSPF route into other OSPF areas. Because the superbackbone appears as another area behind the PE router (acting as ABR), the MP-BGP route derived from the intra-area route is always inserted as an interarea route. The interarea route can then be propagated into other OSPF areas by ABRs within the customer site.

OSPF Superbackbone: Rules

Cisco.com

OSPF superbackbone behaves exactly like Area 0 in regular OSPF:

- PE routers are advertised as Area Border Routers.
- Routes redistributed from BGP into OSPF appear as interarea summary routes or as external routes (based on their original LSA type) in other areas.
- Routes from Area 0 at one site appear as interarea routes in Area 0 at another site.

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-10

Here is a summary of the OSPF superbackbone rules:

- PE routers advertise themselves as ABRs. The superbackbone appears as another area to the CE routers.
- Routes redistributed into MP-BGP from OSPF will appear as interarea routes in other OSPF sites if the original route was an intra-area or interarea route and as external routes if the original route was an external route.

As a consequence of the second rule, routes from the backbone area at one site appear as interarea routes (not as backbone routes) in backbone areas at other sites.

OSPF Superbackbone: Implementation

Cisco.com

- Extended BGP communities are used to propagate OSPF route type across BGP backbone.
- OSPF cost is copied into MED attribute.

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-11

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

A new BGP extended community was defined to carry OSPF route type and OSPF area across the BGP backbone. The format of this community is defined in this table.

BGP Extended Community Format

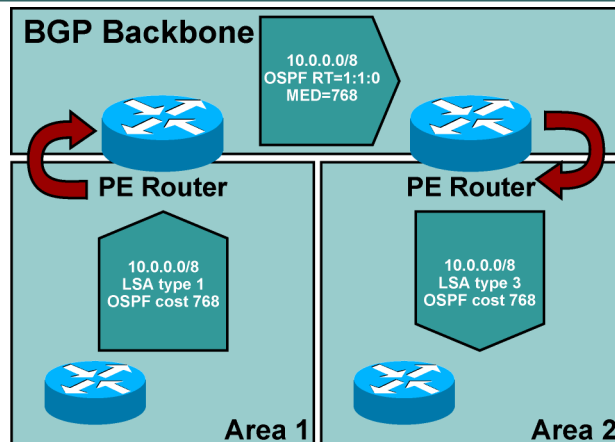
Field	Number of Bytes	Comments
Community type	2	The community type is 0x8000.
OSPF area	4	This field carries the OSPF area from which the route was redistributed into MP-BGP.
LSA type	1	This field carries the OSPF LSA type from which the route was redistributed into MP-BGP.
Option	1	This field is used for external metric type. The low-order bit is set for external 2 routes.

Note The option field in the OSPF route type extended community is not equivalent to the option field in the OSPF LSA.

As in the standard OSPF-BGP redistribution, the OSPF cost is carried in the MED attribute.

OSPF Superbackbone: Implementation (Cont.)

Cisco.com



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

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-12

This figure illustrates the propagation of internal OSPF routes across the MPLS VPN superbackbone.


Example: OSPF Superbackbone Implementation

The sending PE router redistributes the OSPF route into MP-BGP, copies the OSPF cost into the MED attribute, and sets the BGP extended community to indicate the LSA type from which the route was derived.

The receiving PE router redistributes the MP-BGP route back into OSPF and uses the original LSA type and the MED attribute to generate an interarea summary LSA. An interarea summary LSA is always generated because the receiving PE router acts as an ABR between the superbackbone and the OSPF area (or areas).

OSPF Superbackbone: External Routes

Cisco.com



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

© 2004 Cisco Systems, Inc. All rights reserved.

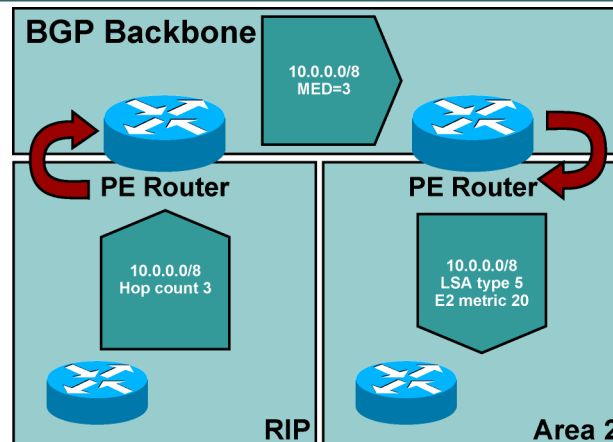
MPLS v2.1—5-13

The external OSPF routes are redistributed into MP-BGP in exactly the same way as the internal OSPF routes. The following describes how the process changes slightly on the receiving PE router:

- For external routes (type 5 LSA), the LSA is reoriginated, with the receiving PE router being the ASBR. The external metric type is copied from the BGP extended community, and the external cost is copied from the MED.
- For NSSA external routes (type 7 LSA), the route is announced to the other OSPF sites as a type 5 LSA external route, because the route has already crossed the area boundary.

OSPF Superbackbone: Mixing Routing Protocols

Cisco.com



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

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-14

The MPLS VPN superbackbone still retains the traditional OSPF-BGP route redistribution behavior for routes that did not originate in OSPF at other sites (and therefore do not carry the OSPF extended BGP community). These routes are inserted into the OSPF topology database as type 5 external routes (or type 7 external routes for NSSA areas), with the default OSPF metric (not the value of MED).

Configuring OSPF PE-CE Routing

This topic describes how to configure a PE-CE OSPF routing session.

Configuring PE-CE OSPF Routing

Cisco.com

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

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

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-15

To configure OSPF as a PE-CE routing protocol, you need to start a separate OSPF process for each VRF in which you want to run OSPF. The per-VRF OSPF process is configured in the same way as a standard OSPF process. You can use all of 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.

Configuring PE-CE OSPF Routing (Cont.)

Cisco.com

```
router(config)#
```

```
router ospf process-id vrf vrf-name  
... Standard OSPF parameters ...
```

- This command starts the per-VRF OSPF routing process.
- The total number of routing processes per router is limited to 32.

```
router(config-router)#
```

```
redistribute bgp as-number subnets
```

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

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-16

OSPF is the only PE-CE routing protocol that is not fully VPN-aware. A separate OSPF process is run for every VRF.

router ospf

To configure an OSPF routing process within a VRF, use the **router ospf** command in global configuration mode. To terminate an OSPF routing process, use the **no** form of this command.

- **router ospf** *process-id* **vrf** *vrf-name*
- **no router ospf** *process-id* **vrf** *vrf-name*

This table describes the parameters for the **router ospf** command.

Syntax Description

Parameter	Description
<i>process-id</i>	Internally used identification parameter for an OSPF routing process. This parameter is locally assigned and can be any positive integer. A unique value is assigned for each OSPF routing process.
<i>vrf-name</i>	Name of the VRF where the OSPF process will reside.

Defaults

No OSPF routing process is defined.

Configuring PE-CE OSPF Routing (Cont.)

Cisco.com

```
router(config)#
```

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

- **OSPF-BGP route redistribution is configured with the redistribute command under the proper address-family command.**
- **Without the OSPF match keyword specified, only internal OSPF routes are redistributed into OSPF.**


© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-17

Use the standard BGP redistribution commands.

Using the OSPF Down Bit

This topic describes how the OSPF down bit is used to address the route loop issue.



Example: OSPF Down Bit

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

This table indicates the process steps that could produce a routing loop.

Process Steps in a Routing Loop

Step	Action
1	The sending PE router receives an intra-area OSPF route.
2	The intra-area OSPF route is redistributed into MP-BGP. An OSPF community is attached to the route to indicate that it was an OSPF route before being redistributed.
3	The receiving PE router redistributes the MP-BGP route into OSPF as an internal interarea summary route.
4	The summary route is propagated across the OSPF area and received by the other PE router attached to the same area.
5	The administrative distance of the OSPF route is better than the administrative distance of the MP-BGP route; therefore, the PE router selects the OSPF route and redistributes the route back into the MP-BGP process, potentially resulting in a routing loop.

OSPF Down Bit: Loop Prevention

Cisco.com

- An additional bit (**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.

© 2004 Cisco Systems, Inc. All rights reserved.

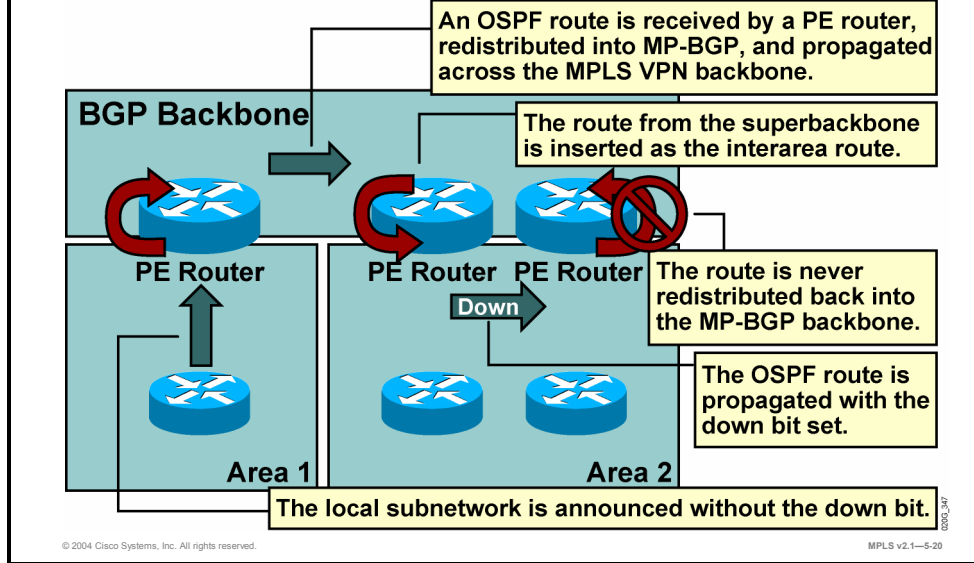
MPLS v2.1—5-19

The following two mechanisms were introduced to prevent route redistribution loops between OSPF (running between PE and CE routers) and MP-BGP running between PE routers:

- One of these mechanisms is the BGP Site of Origin (SOO), which is covered in the “Introducing MPLS VPN Routing Model” lesson of the “MPLS Virtual Private Network Technology” module and detailed further in the “Configuring BGP as the Routing Protocol Between PE and CE Routers” lesson of the “MPLS VPN Implementation” module.
- The other mechanism is the down bit in the options field of the OSPF LSA header.

OSPF Down Bit: Loop Prevention (Cont.)

Cisco.com



The down bit is used between the PE routers to indicate which routes were inserted into the OSPF topology database from the MPLS VPN superbackbone and thus shall not be redistributed back in the MPLS VPN superbackbone. The PE router that redistributes the MP-BGP route as 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.

Example: OSPF Down Bit


The typical usage of the down bit is shown in the figure. The steps that show how the down bit prevents routing loops are detailed in this table.

Process to Prevent Routing Loops

Step	Action
1	The PE router receives an OSPF route.
2	The PE router redistributes the OSPF route into MP-BGP. The MP-BGP route is propagated to the other PE routers.
3	The MP-BGP route is inserted as an interarea route into an OSPF area by the receiving PE router. The receiving PE router sets the down bit in the summary (type 3) LSA.
4	When the other PE routers receive the summary LSA with the down bit set, they do not redistribute the route back into MP-BGP.

Optimizing Packet Forwarding Across the MPLS VPN Backbone

This topic describes how packet forwarding is optimized across the MPLS VPN backbone.



The OSPF superbackbone implementation with MP-BGP has other implications beyond the potential for routing loops between OSPF and BGP.

Example: Optimizing of Packet Forwarding

Consider, for example, the network in the figure. This table indicates a typical flow for routing updates.

Process Steps for Routing Update Flow

Step	Action
1	The PE router redistributes the OSPF route into MP-BGP. The route is propagated to other PE routers as an MP-BGP route. The route is also redistributed into other OSPF areas.
2	The redistributed OSPF route is propagated across the OSPF area with the down bit set.
3	The ingress PE router receives an MP-IBGP route with an administrative distance of 200 and an OSPF route with an administrative distance of 110. The OSPF route is preferred over the MP-IBGP route, and the data packets flow across customer sites, not directly over the MPLS VPN backbone.

Optimizing of Packet Forwarding Across the MPLS VPN Backbone (Cont.)

Cisco.com

- The PE routers ignore OSPF routes with the down bit set for routing purposes:
 - These routes originated at other sites; therefore, the traffic toward them should go via the MP-BGP backbone.
- The **routing** bit is not set on OSPF routes with the down bit set:
 - These routes do not enter the IP routing table, even when they are selected as the best routes using the SPF algorithm.

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-22


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

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

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

Optimizing of Packet Forwarding Across the MPLS VPN Backbone (Cont.)

Cisco.com



© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-23

With the new route OSPF selection rules in place, the packet forwarding in the network shown in the figure follows the desired path. The process steps are described in this table.

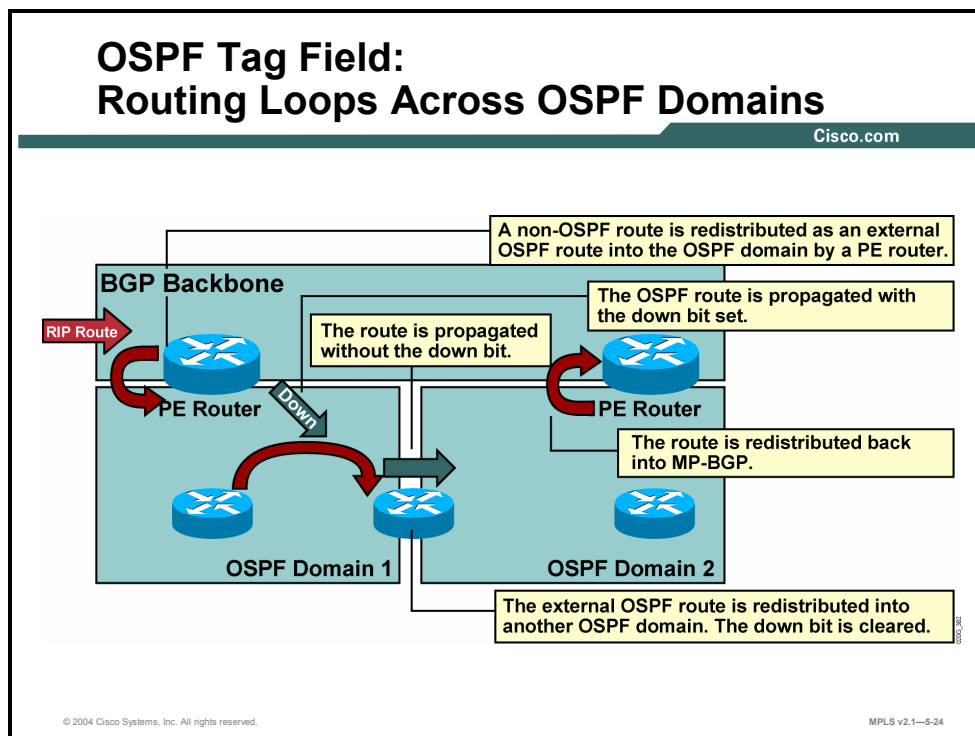
Process Steps for Optimizing Packet Forwarding

Step	Action
1	The OSPF route is redistributed into MP-BGP by a PE router and propagated to other PE routers.
2	The receiving PE routers redistribute the MP-BGP route into OSPF.

Other PE routers might receive the MP-BGP and OSPF routes but will ignore the OSPF route for routing purposes because it has the down bit set. The data packets will flow across the MPLS VPN backbone, following only the MP-BGP routes, not the OSPF routes derived from the MP-BGP routes.

Using the OSPF Tag Field

This topic describes how the OSPF tag field is used to address the root loop issue.



The down bit stops the routing loops between MP-BGP and OSPF. The down bit cannot, however, stop the routing loops when redistribution between multiple OSPF domains is involved.

Example: OSPF Tag Field

The figure illustrates this example. The routing loop in this network occurs as part of the steps outlined in this table.

Process Steps for Routing Loops Across OSPF Domains

Step	Action
1	The PE router redistributes a non-OSPF route into an OSPF domain as an external route. The down bit is set because the route should not be redistributed back into MP-BGP.
2	A CE router redistributes the OSPF route into another OSPF domain. The down bit is lost if the CE router does not understand this OSPF extension.
3	The OSPF route is propagated through the other OSPF domain with the down bit cleared.
4	A PE-router receives the OSPF route; the down bit is not set, so the route is redistributed back into the MP-BGP backbone, resulting in a routing loop.

OSPF Tag Field: Operation

Cisco.com

- **The tag field in external OSPF routes is used to detect cross-domain routing loops.**
- **PE routers set the tag field to the BGP AS number when redistributing non-OSPF routes from MP-BGP into OSPF.**
- **The tag field is propagated between OSPF domains when the external OSPF routes are redistributed between OSPF domains.**
- **PE routers filter external OSPF routes to MP-BGP with OSPF tag field AS numbers matching BGP AS numbers.**

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-25

The routing loops introduced by route redistribution between OSPF domains can be solved with the help of the tag field, using standard OSPF-BGP redistribution rules.

In standard OSPF-BGP or OSPF-OSPF redistribution, the following rules apply:

- Whenever a router redistributes a BGP route into OSPF, the tag field in the type 5 (or type 7) LSA is set to the AS number of the redistributing router.
- The tag field from an external OSPF route is propagated across OSPF domains when the external OSPF route is redistributed into another OSPF domain.
- In addition to these standard mechanisms, PE routers filter external OSPF routes based on their tag field and do not redistribute, into MP-BGP, routes with a tag field equal to the BGP AS number.

OSPF Tag Field: Usage Guidelines

Cisco.com

- Internal OSPF routes have no tag field.
- This technique does not detect cross-domain routing information loops for routes inserted as internal OSPF routes by the PE routers.
- The tag field can be set manually on the router, redistributing routes between OSPF domains with the `redistribute ospf source-process-id tag value` command.
- Alternatively, only the internal OSPF routes can be redistributed into MP-BGP on the PE routers.

© 2004 Cisco Systems, Inc. All rights reserved.

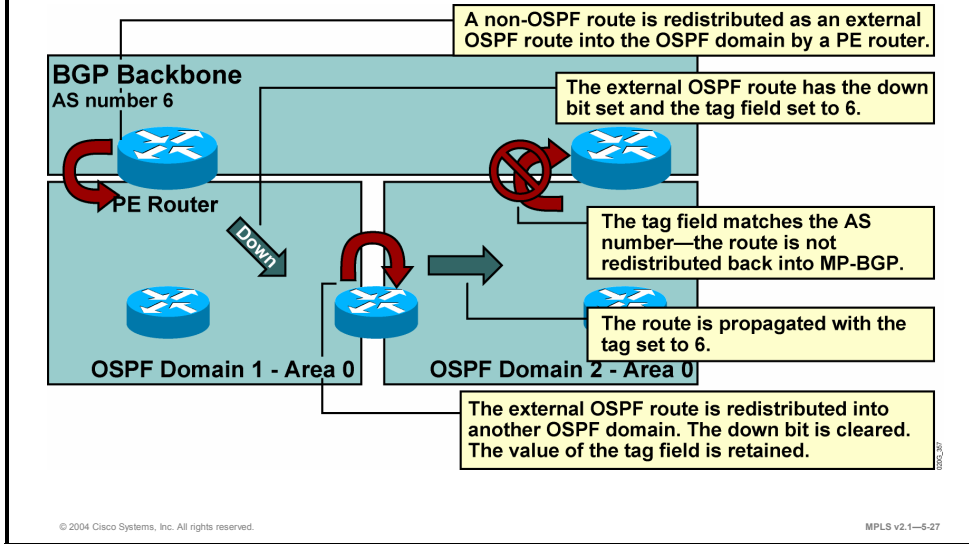
MPLS v2.1—5-26

The OSPF tag field is present only in the external OSPF routes (type 5 LSA or type 7 LSA). This technique, therefore, cannot detect cross-domain loops involving internal OSPF routes. Here are the two manual methods that you can use to overcome this OSPF limitation:

- You can set the tag field manually on the router, redistributing routes between OSPF domains using the `redistribute ospf source-process-id tag value` command.
- The PE router can be configured to redistribute only internal OSPF routes into MP-BGP.

OSPF Tag Field: Routing Loop Prevention

Cisco.com



The OSPF tag field can be used to prevent routing loops when the redistribution is done between OSPF domains.

Example: OSPF Tag Field—Routing Loop Prevention

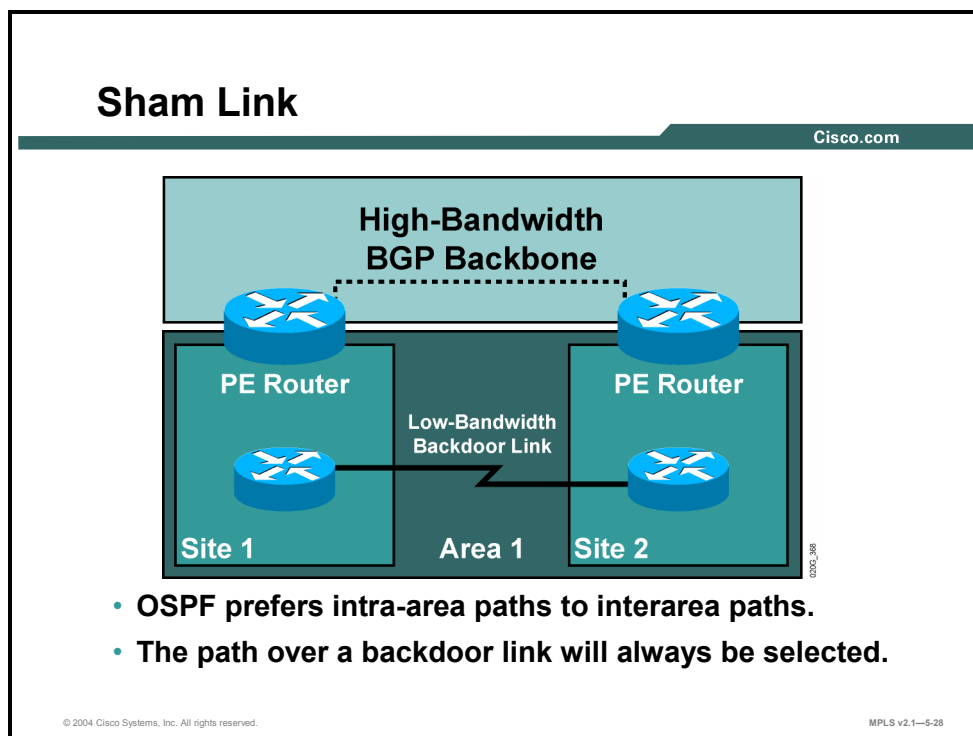
The figure illustrates this example. This table lists the steps in this process.

Process Steps to Prevent Routing Loops

Step	Action
1	A non-OSPF route is redistributed as an external OSPF route by a PE router. The tag field is set to the BGP AS number, and the down bit is set.
2	The redistributed route is propagated across the OSPF domain.
3	When the route is redistributed into another OSPF domain, the tag field is propagated, but the down bit is cleared.
4	Another PE router receives the external OSPF route and filters the route based on the tag field. The route is not redistributed into MP-BGP.

What Is a Sham Link?

This topic describes the features of a sham link.



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.

Example: Sham Link

The figure illustrates the 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.

Sham Link (Cont.)

Cisco.com

- **A logical intra-area link.**
- **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.**

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-29

If the backdoor links between sites are used only for backup purposes and do not participate in the VPN service, the default route selection shown in the preceding figure is not acceptable. To reestablish the desired path selection over the MPLS VPN backbone, you must create an additional OSPF intra-area (logical) link between ingress and egress VRFs on the relevant PE routers. This link is called a sham link.

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

Sham Link (Cont.)

Cisco.com

When a sham-link route is preferred by OSPF:

- The OSPF route is not redistributed to MP-BGP.
- Instead, the router on the other end of the sham link performs the redistribution.
- The forwarding information from the MP-BGP route is used.

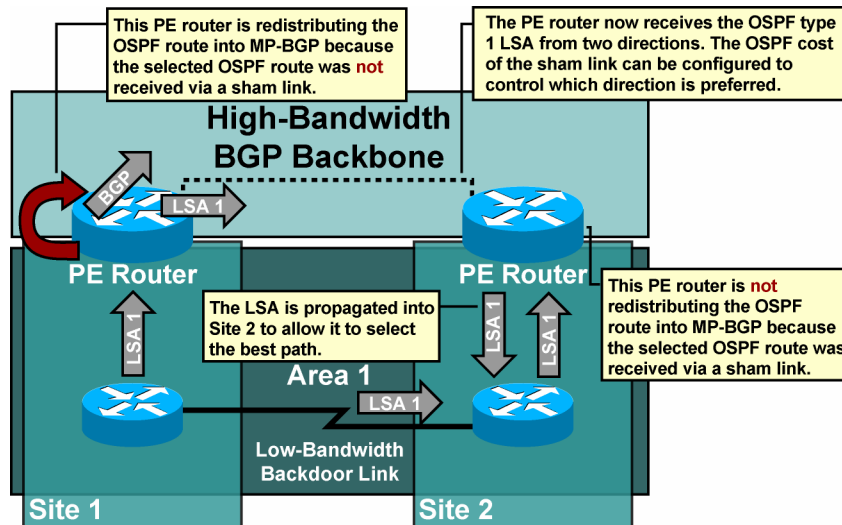
© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-30

A cost is configured with each sham link. This cost is used to decide whether traffic will be sent over the backdoor path or the sham-link path. When a sham link is configured between PE routers, the PE routers can populate the VRF routing table with the OSPF routes learned over the sham link.

Sham Link (Cont.)

Cisco.com



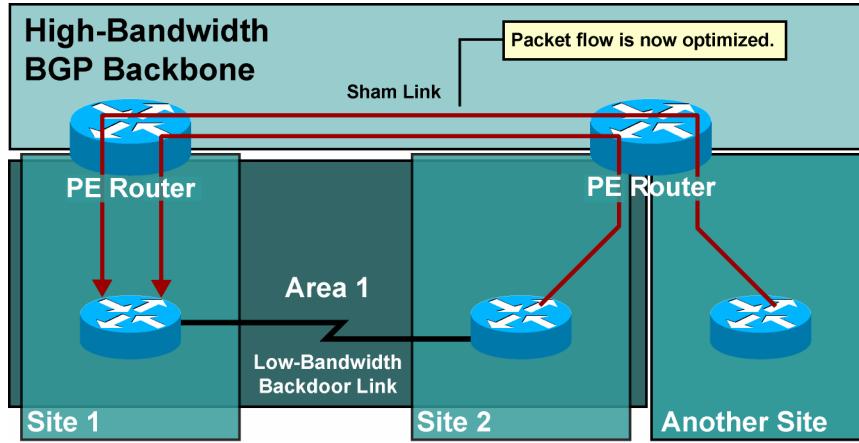
© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-31

Because the sham link is seen as an intra-area link between PE routers, an OSPF adjacency is created and database exchange (for the particular OSPF process) occurs across the link. The PE router can then flood LSAs between sites from across the MPLS VPN backbone. As a result, the desired intra-area connectivity is created.

Sham Link (Cont.)

Cisco.com



© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-32

The implementation results in optimum packet flow.

Configuring a Sham Link

This topic describes how to configure a sham link.

Configuring a Sham Link

Cisco.com

- **A separate /32 address space is required in each PE router for the 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 shamlink**
 - **Must belong to the VRF**
 - **Must not be advertised by OSPF**
 - **Must be advertised by BGP**

© 2004 Cisco Systems, Inc. All rights reserved.MPLS v2.1—5-33

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

The following 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

Configuring a Sham Link (Cont.)

Cisco.com

```
router(config-router)#
```

```
area area-id sham-link source-address destination-address cost number
```

- This command was introduced in Cisco IOS Release 12.2(8) T.
- The sham link belongs to the specified area.
- Sham-link packets sent across the MPLS VPN backbone will have the specified source and destination addresses.
- When the SPF algorithm is executed, the sham link will have the specified cost.

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-34

To configure a sham-link interface on a PE router in an MPLS VPN backbone, use the **area sham-link cost** command in global configuration mode. To remove the sham link, use the **no** form of this command.

- **area** *area-id* **sham-link** *source-address destination-address* **cost** *number*
- **no area** *area-id* **sham-link** *source-address destination-address* **cost** *number*

This table describes the parameters for the **area sham-link cost** command.

Syntax Description

Parameter	Description
<i>area-id</i>	ID number of the OSPF area assigned to the sham link. Valid values: numeric value or valid IP address. There is no default.
<i>source-address</i>	IP address of the source PE router in the format: <i>ip-address [mask]</i> .
<i>destination-address</i>	IP address of the destination PE router in the format: <i>ip-address [mask]</i> .
<i>number</i>	OSPF cost to send IP packets over the sham-link interface. Valid values are from 1 to 65535.

Defaults


There is no default behavior or values.

Command Modes

Use this command in global configuration mode.

Sample Sham-Link Configuration

Cisco.com



```
ip vrf Customer_A
rd 115:43
route-target both 115:43
!
interface Loopback11
ip forwarding vrf Customer_A
ip address 10.2.1.1 255.255.255.255
!
interface serial 1/0/1
ip forwarding vrf Customer_A
ip address 10.1.0.1 255.255.255.252
!
router ospf 11 vrf Customer_A
network 10.1.0.1 0.0.0.3 area 1
redistribute bgp 3 subnets
area 1 sham-link 10.2.1.1 10.2.1.2 cost 10
!
router bgp 3
address-family ipv4 vrf Customer_A
network 10.2.1.1 mask 255.255.255.255
redistribute ospf 11 match internal
```

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-35

A sham link is used only to affect the OSPF intra-area path selection of the PE and CE routers.

Example: Sample Sham-Link Configuration

The figure illustrates this example. The PE router also uses the information received from MP-BGP to set the outgoing label stack of incoming packets and to decide to which egress PE router to label-switch the packets.

The figure shows a sample MPLS VPN topology in which a sham-link configuration is necessary. A VPN client has two sites connected by a backdoor link. A sham link has been configured between the two PE routers.

Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- **OSPF divides a network into areas. Individual areas connect to a common backbone in a two-tier hierarchical model.**
- **The MPLS VPN routing model introduces a BGP backbone into the OSPF network of the customer. OSPF is run at each site, while MP-BGP is used to propagate routes between each site.**
- **MPLS VPN architecture implements the MP-BGP backbone as a new OSPF superbackbone above existing areas. This new superbackbone is transparent to the existing OSPF network.**
- **OSPF PE-CE routing is implemented as a new routing process. One routing process per VRF is configured using the VRF name.**
- **The OSPF down bit prevents routing loops by denying OSPF routes originating from MP-BGP with the down bit enabled from being redistributed back into MP-BGP.**

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-36

Summary (Cont.)

Cisco.com

- **PE routers have to ignore all OSPF routes with the down bit set. This is because these routes originated in the MP-BGP backbone and the MP-BGP route should be used as the optimum route toward the destination.**
- **The OSPF tag field prevents routing loops of external routes between different OSPF domains. The tag field is set to the AS number of the originating MP-BGP router to deny the packets return.**
- **A sham link is required between any two VPN sites that belong to the same OSPF area and share an OSPF backdoor link.**
- **Use the area sham-link cost command to configure the sham link.**

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-37

Configuring BGP as the Routing Protocol Between PE and CE Routers

Overview

This lesson explains the PE-CE routing protocol configuration steps required when you are using BGP as the routing protocol between PE and CE routers.

It is important to understand not only what you can configure between PE and CE routers when you are setting up MPLS VPNs, but also how to accomplish the configuration successfully. This lesson looks at the configuration parameters that you need to configure an MPLS VPN PE-CE routing exchange.

Objectives

Upon completing this lesson, you will be able to describe how to configure BGP as the routing protocol between CE and PE routers. This ability includes being able to meet these objectives:

- Describe how to configure a per-VRF BGP routing context
- Explain the reason for limiting the number of routes in a VRF
- Describe how to limit the number of prefixes received from a BGP neighbor
- Describe how to limit the total number of VRF routes
- Identify the issues encountered when a customer wants to reuse the same AS number on several sites
- Identify the issues encountered when a customer site links two VPNs
- Describe how to implement SOO for loop prevention

Configuring a Per-VRF BGP Routing Context

This topic describes how to configure a per-VRF BGP routing context.

Configuring per-VRF BGP Routing Context

Cisco.com

```
Router(config)#  
router bgp as-number  
  address-family ipv4 vrf vrf-name  
    ... Per-VRF BGP definitions ...
```

- **Select per-VRF BGP context with the address-family command.**
- **Configure CE EBGP neighbors in the VRF context, not in the global BGP configuration.**
- **CE neighbors have to be activated with the neighbor activate command.**

© 2004 Cisco Systems, Inc. All rights reserved. MPLS v2.1-5.3

Select the VRF routing context with the **address-family ipv4 vrf *vrf-name*** command in the BGP routing processes. All per-VRF routing protocol parameters (network numbers, passive interfaces, neighbors, filters, and so on) are configured under this address family.

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

Note You always have to configure a BGP address family for each VRF and configure route redistribution into BGP for each VRF, even if you do not use BGP as the PE-CE routing protocol.

Several BGP options have different default values when you configure the per-VRF BGP routing context, as follows:

- BGP synchronization is disabled (default = enabled).
- Autosummarization (automatic generation of classful networks out of subnetworks redistributed into BGP) is disabled (default = enabled). This is because the MPLS VPN backbone has to propagate customer subnetworks unchanged to facilitate transparent end-to-end routing between customer sites. Redistribution of internal BGP routes into IGP is enabled (default = disabled).

Note The common parameters defined in router configuration mode are inherited by all address families defined for this routing process and can be overridden for each individual address family.

address-family ipv4

To enter address family configuration mode for configuring routing sessions (such as BGP) that use standard IPv4 address prefixes, use the **address-family ipv4** command in router configuration mode. To disable address family configuration mode, use the **no** form of this command.

- **address-family ipv4** [**multicast** | **unicast** | **vrf** *vrf-name*]
- **no address-family ipv4** [**multicast** | **unicast** | **vrf** *vrf-name*]

This table describes the parameters for the **address-family ipv4** command.

Syntax Description

Parameter	Description
multicast	(Optional) Specifies IPv4 multicast address prefixes.
unicast	(Optional) Specifies IPv4 unicast address prefixes.
vrf <i>vrf-name</i>	(Optional) Specifies the name of the VRF instance to associate with subsequent IPv4 address family configuration mode commands.

Defaults

IPv4 address prefixes are not enabled. Unicast address prefixes are the default when IPv4 address prefixes are configured.


Command Modes

Use this command in router configuration mode.

Configuring per-VRF BGP Routing Context (Cont.)

Cisco.com

```
router bgp 65001
neighbor 10.200.1.2 remote-as 115
network 10.1.0.0 mask 255.255.0.0
```



```
ip vrf Site_A
rd 115:317
route-target both 115:317
```

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1-5-4

The PE router can be defined as a BGP neighbor.

Example: Configuring per-VRF BGP Routing Context

The figure shows BGP that is activated on the CE router, and the PE router is defined as a BGP neighbor. Similarly, the CE router is defined as a BGP neighbor and activated under the **address-family ipv4 vrf Customer_A** command.

What Are the Reasons for Limiting the Number of Routes in a VRF?

This topic explains the reason for limiting the number of routes in a VRF.

Limiting the Number of Routes in a VRF

Cisco.com

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

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-5

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

To limit the effect of configuration errors and malicious user behavior, Cisco IOS software offers the following two features that limit the number of routes and resource consumption that a VPN user can take advantage of at a PE router:

- The BGP maximum-prefix feature limits the number of routes that an individual BGP peer can send.
- The VRF route limit 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 MP-IBGP.

Limiting the Number of Prefixes Received from a BGP Neighbor

This topic describes how to limit the number of prefixes received from a BGP neighbor.

Limiting the Number of Prefixes Received from a BGP Neighbor

Cisco.com

```
Router (config-router-af) #  
neighbor ip-address maximum-prefix maximum [threshold]  
[warning-only]
```

- Controls how many prefixes can be received from a neighbor
- Optional *threshold* parameter specifies the percentage where a warning message is logged (default is 75 percent)
- Optional **warning-only** keyword specifies the action on exceeding the maximum number (default is to drop peering)

© 2004 Cisco Systems, Inc. All rights reserved. MPLS v2.1—6-6

neighbor maximum-prefix

To control how many prefixes can be received from a neighbor, use the **neighbor maximum-prefix** command in router configuration mode. To disable this function, use the **no** form of this command.

- **neighbor** {*ip-address* | *peer-group-name*} **maximum-prefix** *maximum* [*threshold*] [**warning-only**]
- **no neighbor** {*ip-address* | *peer-group-name*} **maximum-prefix** *maximum* [*threshold*] [**warning-only**]

This table describes the parameters for the **neighbor maximum-prefix** command.

Syntax Description

Parameter	Description
<i>ip-address</i>	IP address of the neighbor.
<i>peer-group-name</i>	Name of a BGP peer group.
<i>maximum</i>	Maximum number of prefixes allowed from this neighbor.
<i>threshold</i>	(Optional) Integer identifying at what percentage of the maximum the router starts to generate a warning message. The range is 1 to 100; the default is 75 (percent).
warning-only	(Optional) Allows the router to generate a log message when the maximum is exceeded instead of terminating the peering.

Defaults

Default is disabled; there is no limit on the number of prefixes.

Limiting the Total Number of VRF Routes

This topic describes how to limit VRF routes.

Limiting the Total Number of VRF Routes

Cisco.com

- **The VRF maximum routes *limit* command limits the number of routes that are imported into a VRF:**
 - Routes coming from CE routers
 - Routes coming from other PE routers (imported routes)
- **The route limit is configured for each VRF.**
- **If the number of routes exceeds the route limit:**
 - A syslog message is generated.
 - The Cisco IOS software can be configured to reject routes (optional).

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-7

The VRF route limit, contrary to the BGP maximum-prefix limit, limits the overall number of routes in a VRF regardless of their origin. Similar to the BGP maximum-prefix feature, the network operator might be warned via a syslog message when the number of routes exceeds a certain threshold. Additionally, you can configure Cisco IOS software to ignore new VRF routes when the total number of routes exceeds the maximum configured limit.

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

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

Limiting the Total Number of VRF Routes (Cont.)

Cisco.com

```
Router(config-vrf)#
```

```
maximum routes limit {warn threshold | warn-only}
```

- This command configures the maximum number of routes accepted into a VRF:
 - The *limit* parameter is the route limit for the VRF.
 - The *warn threshold* parameter is the percentage value over which a warning message is sent to syslog.
 - The *warn-only* parameter the PE continues accepting routes after the configured limit.
- Syslog messages generated by this command are rate-limited.

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-8

maximum routes

To limit the maximum number of routes in a VRF to prevent a PE router from importing too many routes, use the **maximum routes** command in VRF configuration submode. To remove the limit on the maximum number of routes allowed, use the **no** form of this command.

- **maximum routes** *limit* {*warn threshold* | **warn-only**}
- **no maximum routes**

This table describes the parameters for the **maximum routes** command.

Syntax Description

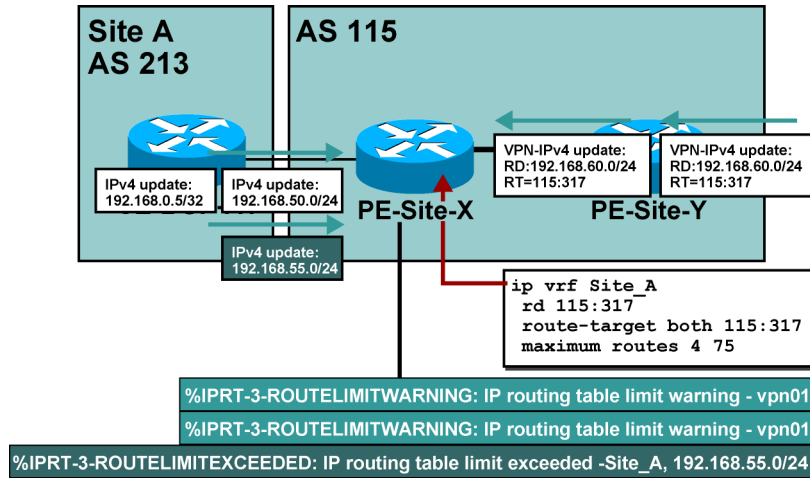
Parameter	Description
<i>limit</i>	Specifies the maximum number of routes allowed in a VRF. You may select from 1 to 4,294,967,295 routes to be allowed in a VRF.
<i>warn threshold</i>	Rejects routes when the threshold limit is reached. The threshold limit is a percentage of the limit specified, from 1 to 100.
warn-only	Issues a syslog error message when the maximum number of routes allowed for a VRF exceeds the threshold. However, additional routes are still allowed.

Defaults

This command has no default behavior or values.

Limiting the Total Number of VRF Routes (Cont.)

Cisco.com



© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1-5-9

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

Example: Limiting the Total Number of VRF Routes

In this figure, the network designer has decided to limit the number of routes in a VRF to 4, with the warning threshold being set at 75 percent (or 3 routes).

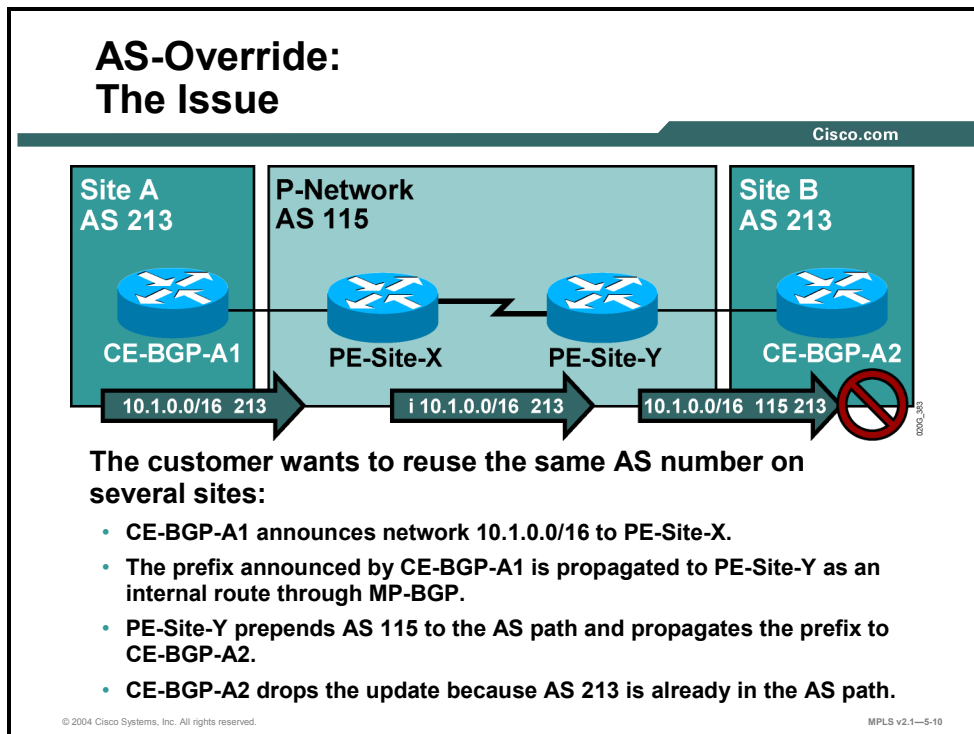
When the first two routes are received and inserted into the VRF, the router accepts them. When the third route is received, a warning message is generated, and the message is repeated with the insertion of the fourth route.

Note The syslog messages are rate-limited to prevent indirect denial-of-service attacks on the network management station.

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

Identifying AS-Override Issues

This topic identifies the issues encountered when a customer wants to reuse the same AS number on several sites.



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

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

AS-Override: Implementation

Cisco.com

- **New AS path update procedures have been implemented to reuse the same AS number on all VPN sites.**
- **The procedures allow the use of private and public AS numbers.**
- **The same AS number may be used for all sites.**

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-11

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

AS-Override: Implementation (Cont.)

Cisco.com

With AS-override configured, the AS path update procedure on the PE router is as follows:

- **If the first AS number in the AS path is equal to the neighboring AS, it is replaced with the provider AS number.**
- **If the first AS number has multiple occurrences (because of AS path prepend), all occurrences are replaced with the provider AS number.**
- **After this operation, the provider AS number is prepended to the AS path.**

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-12

The modified AS path update procedure is called AS-override, which is described here:

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

AS-Override: Command

Cisco.com

```
Router(config-router-af)#
```

```
neighbor ip-address as-override
```

- This command configures the AS-override AS path update procedure for the specified neighbor.
- AS-override is configured for CE EBGP neighbors in the VRF address family of the BGP process.

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-13

neighbor as-override

To configure a PE router to override a site AS number with a provider AS number, use the **neighbor as-override** command in router configuration mode. To remove VPNv4 prefixes from a specified router, use the **no** form of this command.

- **neighbor ip-address as-override**
- **no neighbor ip-address as-override**

This table describes the parameters for the **neighbor as-override** command.

Syntax Description


Parameter	Description
<i>ip-address</i>	Specifies the router IP address to override with the AS number provided.

Defaults

This command has no default behavior or values.

Example: AS-Override

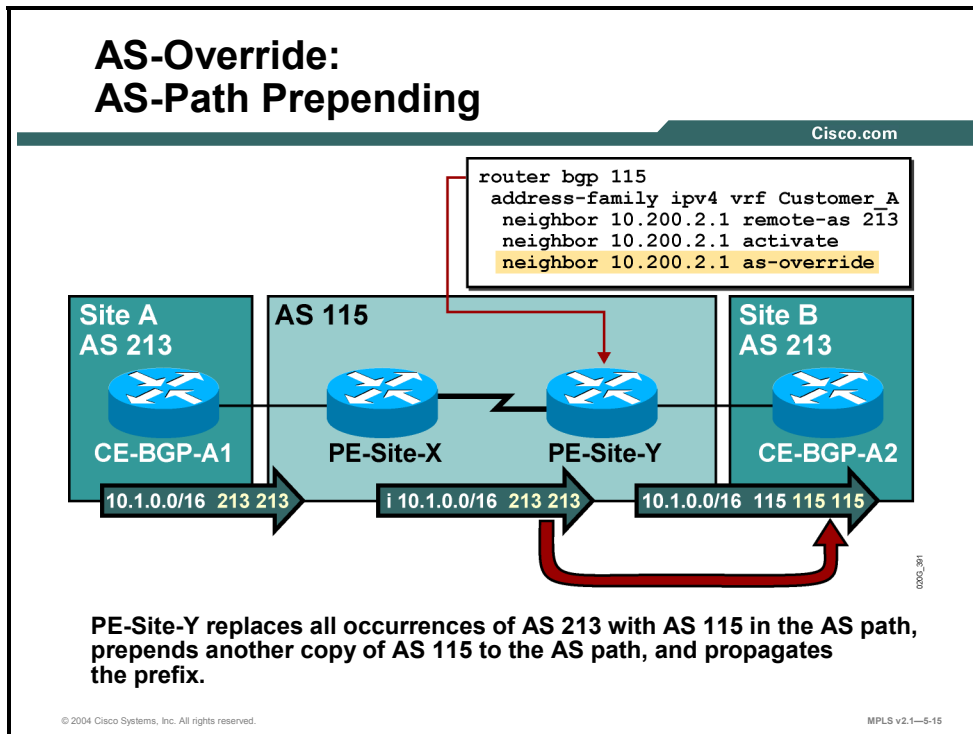
In this figure, customer sites A and B use BGP to communicate with the MPLS VPN backbone. Both sites use AS 213. Site B would drop the update sent by site A without the AS-override mechanism.



The AS-override mechanism, configured on the PE-Site-Y router, replaces the customer AS number (213) with the provider AS number (115) before sending the update to the customer site. An extra copy of the provider AS number is prepended to the AS path during the standard EBGp update process.

Example: AS-Path Prepending

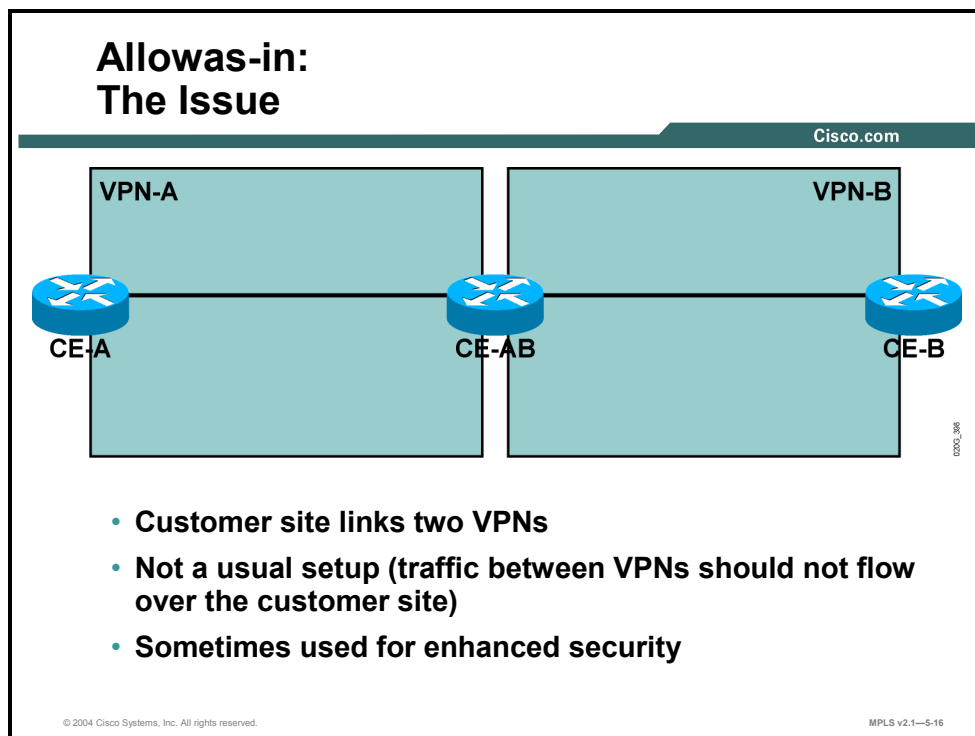
In this figure, the customer is using AS prepending to influence BGP path selection within the MPLS VPN backbone.



The PE router has to send a route with an AS path containing multiple copies of the customer AS number. In this case, all the leading copies of the customer AS number are replaced with the provider AS number (resulting in two occurrences of the provider AS number in the example), and the third occurrence of the provider AS number is prepended to the BGP update before it is sent to the CE router.

Identifying Allowas-in Issues

This topic identifies the issues encountered when a customer site links two VPNs.

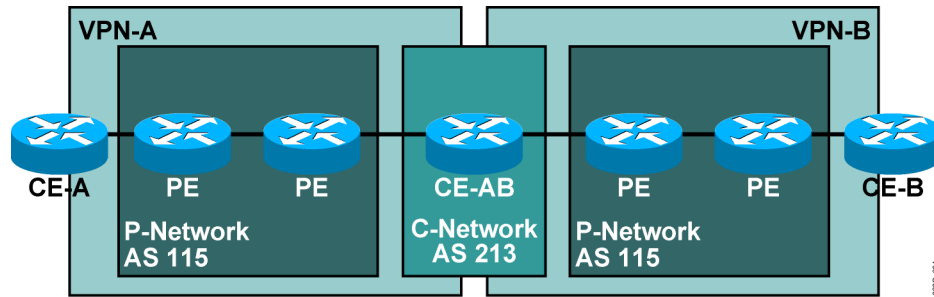


In some security-conscious implementations, customer VPNs are linked by a customer router that performs security functions, such as access filtering or access logging.

Note This setup is not usual because it deviates from the basic goal of MPLS VPN—replacing the hub-and-spoke routing of a traditional overlay VPN with optimum any-to-any routing.

Allowas-in: The Issue (Cont.)

Cisco.com



- **VPN perspective:** VPN-A is connected to VPN-B via CE-BGP-A1.
- **Physical topology:** The CE router is connected to PE routers.
- **MPLS VPN perspective:** The CE router has two links into the P-network.
- **BGP perspective:** The CE router has two connections to AS 115.

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-17


The setup in which a customer router links two VPNs in an MPLS VPN backbone can be viewed from several different perspectives, as follows:

- From the VPN perspective, a CE router links two VPNs.
- From the physical perspective, the CE router is connected through two separate links (physical or logical interface) to one or two PE routers.
- In MPLS VPN terms, the CE router has two links into the P-network.

There is no problem with the proposed customer setup if the setup is analyzed through these perspectives. All of the potential setups represent valid connectivity or routing options. The problem occurs when the setup is analyzed through the BGP perspective, in which the CE router has to propagate routes between two PE routers, which are both in the same AS.

Allowas-in: The Issue (Cont.)

Cisco.com



- PE-1 announces network 10.1.0.0/16 to CE-BGP-A1.
- CE-BGP-A1 prepends its AS number to the AS path and propagates the prefix to PE-2.
- PE-2 drops the update because its AS number is already in the AS path. AS-override is needed on CE-BGP-A1, but that would require a Cisco IOS software upgrade on the CE router.

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-18

Example: Allowas-in

This example is similar to the situation in which two customer sites use the same AS number.

Allowas-in: Implementation

Cisco.com

The allowas-in BGP option disables the AS path check on the PE router:

- The number of occurrences of the PE router AS number is limited to suppress real routing loops.
- The limit has to be configured.
- The PE router will **reject** the update only if its AS number appears in the AS path more often than the configured limit.

© 2004 Cisco Systems, Inc. All rights reserved. MPLS v2.1—5-19

The BGP loop prevention rules prevent a PE router from accepting the routing update sent by the CE router if that routing update already contains the AS number of the MPLS VPN backbone (which it will if the CE router is propagating routes between two VPNs).

The solution to this BGP routing problem is that AS-override has to be used on the CE router. This solution requires a very recent version of Cisco IOS software (Cisco IOS Release 12.0 T or later) on the CE router. The solution is not enforceable in every customer situation.

Networks need to support topologies in which a CE router with no AS-override support links two VPNs. A specific need exists to modify the BGP loop prevention mechanism on the PE routers. The allowas-in feature supports situations in which the PE router receives routes with its own AS number already in the AS path.

With this feature configured on a BGP neighbor of the PE router, the PE router would not drop incoming BGP updates with its AS number in the AS path if the updates are received from that neighbor. To prevent real BGP routing information loops, the number of occurrences of the MPLS VPN backbone AS number can be limited and incoming updates that exceed the limit can be dropped.

Allows-in: Command

Cisco.com

```
Router (config-router) #
```

```
neighbor allows-in number
```

- This command disables the traditional BGP AS path check.
- An incoming update is rejected only if the AS number of the PE router appears in the AS path more often than the configured limit.

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-20

neighbor allows-in

To configure PE routers to allow readvertisement of all prefixes containing duplicate AS numbers, use the **neighbor allows-in** command in router configuration mode. To disable readvertisement of the AS number of a PE router, use the **no** form of this command.

- **neighbor allows-in** *number*
- **no neighbor allows-in** *number*

This table describes the parameters for the **neighbor allows-in** command.

Syntax Description

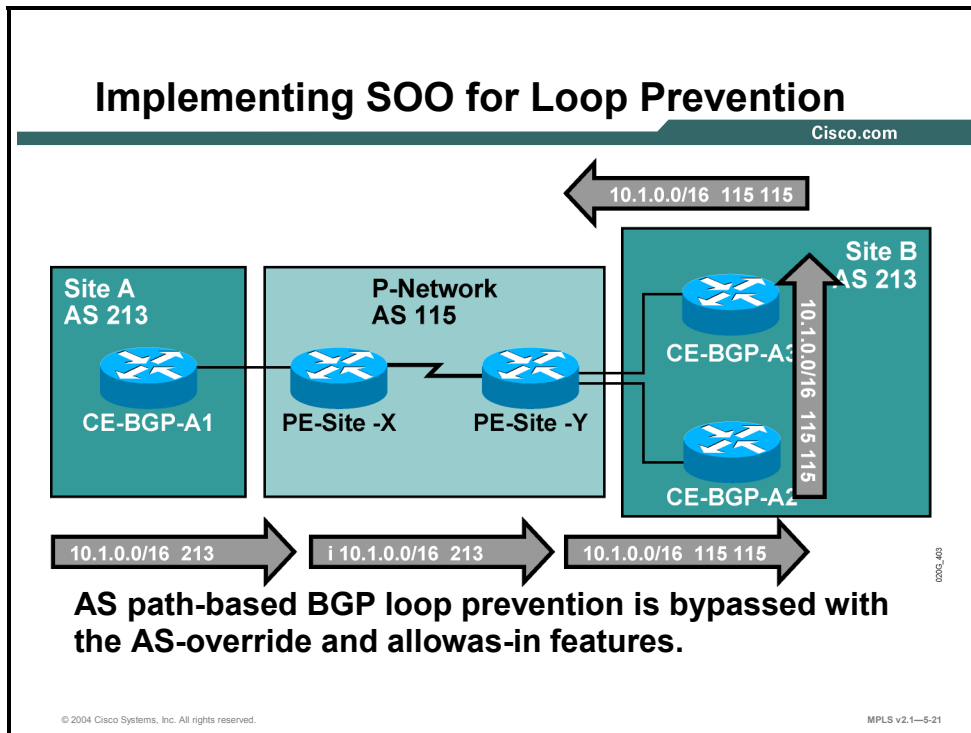
Parameter	Description
<i>number</i>	Specifies the number of times to allow advertisement of the AS number of a PE router. Valid values are from 1 to 10 times.

Defaults

This command has no default behavior or values.

Implementing SOO for Loop Prevention

This topic describes how to implement SOO for loop prevention.



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

The ability to still detect loops can present a particular problem when BGP is mixed with other PE-CE routing protocols. The SOO extended BGP community can be used as an additional loop prevention mechanism in these scenarios.

Note SOO and any other loop prevention mechanisms are needed only for customer networks with multihomed sites. Loops can never occur in customer networks that have only stub sites.

Implementing SOO for Loop Prevention (Cont.)

Cisco.com

- **The SOO attribute (extended BGP community) can be used to prevent loops in these scenarios.**
- **The SOO attribute is needed only for multihomed sites.**
- **When EBGP is run between PE and CE routers, the SOO attribute is configured through a route map command.**
- **For other routing protocols, the SOO attribute can be applied to routes learned through a particular VRF interface during the redistribution into BGP.**

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-22

Here are the two ways to set the SOO attribute on a BGP route:

- For routes received from BGP-speaking CE routers, the SOO attribute is configured by the incoming route map on the PE router.
- For all other routes, a route map setting the SOO attribute is applied to the incoming interface. The SOO attribute, as set by the route map, is attached to the BGP route when an IGP route received through that interface is redistributed into BGP.

Outgoing filters based on the SOO attribute also depend on the routing protocol used, as described here:

- Where EBGP is used as the PE-CE routing protocol, outbound route maps can be used on the PE router to deny routes matching particular SOO values.
- For all other routing protocols, filtering is performed on the basis of the SOO route map configured on the outgoing interface before the update is sent across that interface to the CE router.

Implementing SOO for Loop Prevention (Cont.)

Cisco.com

Inbound EGBP Update

Router(config)#

```
route-map name permit seq
  match conditions
  set extcommunity soo extended-community-value
```

- **Creates a route map that sets the SOO attribute.**

Router(config-router-af)#

```
neighbor ip-address route-map name in
```

- **Applies an inbound route map to the CE EGBP neighbor.**

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-23

set extcommunity

To set the extended communities attribute, use the **set extcommunity** command in route map configuration mode. To delete the entry, use the **no** form of this command.

- **set extcommunity** {**rt** *extended-community-value* [**additive**] | **soo** *extended-community-value*}
- **no set extcommunity**
- **set extcommunity** *extcommunity-type community-number* [**additive**]
- **no set extcommunity** *extcommunity-type community-number* [**additive**]

This table describes the parameters for the **set extcommunity** command.

Syntax Description

Parameter	Description
rt	Specifies the RT extended community attribute.
soo	Specifies the SOO extended community attribute.
<i>extended-community-value</i>	Specifies the value to be set. The value can be one of the following combinations: <ul style="list-style-type: none">■ <i>autonomous-system-number:network-number</i>■ <i>ip-address:network-number</i> The colon is used to separate the AS number from the network number or the IP address from the network number.
additive	(Optional) Adds space after the closing parenthesis. Adds the extended community to the already existing extended communities.

Defaults

No BGP extended community attributes are set by the route map.

neighbor route-map

To apply a route map to incoming or outgoing routes, use the **neighbor route-map** command in address family or router configuration mode. To remove a route map, use the **no** form of this command.

- **neighbor** {*ip-address* | *peer-group-name*} **route-map** *map-name* {**in** | **out**}
- **no neighbor** {*ip-address* | *peer-group-name*} **route-map** *map-name* {**in** | **out**}

This table describes the parameters for the **neighbor route-map** command.

Syntax Description

Parameter	Description
<i>ip-address</i>	IP address of the neighbor.
<i>peer-group-name</i>	Name of a BGP or MP-BGP peer group.
<i>map-name</i>	Name of a route map.
in	Applies route map to incoming routes.
out	Applies route map to outgoing routes.

Implementing SOO for Loop Prevention (Cont.)

Cisco.com

Other Inbound Routing Updates

Router(config)#

```
route-map name permit seq
  match conditions
  set extcommunity soo extended-community-value
```

- **Creates a route map that sets the SOO attribute.**

Router(config-if)#

```
ip vrf sitemap route-map-name
```

- **Applies a route map that sets SOO extended community attribute to inbound routing updates received from this interface.**

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-24

ip vrf sitemap

To set the SOO extended community attribute, use the **ip vrf sitemap** command in interface configuration mode. To delete the entry, use the **no** form of this command.

- **ip vrf sitemap** *route-map-name*
- **no ip vrf sitemap** *route-map-name*

This table describes the parameters for the **ip vrf sitemap** command.

Syntax Description

Parameter	Description
<i>route-map-name</i>	Sets the name of the route map to be used.

Defaults

No route map is used to set the SOO extended community attribute.

Implementing SOO for Loop Prevention (Cont.)

Cisco.com

```
Router(config)#
```

```
ip extcommunity-list number permit soo value
!
route-map name deny seq
  match extcommunity number
!
route-map name permit 9999
```

- Defines a route map that discards routes with the desired SOO value.

```
Router(config-router-af)#
```

```
neighbor ip-address route-map name out
```

- Applies the route map to outbound updates sent to the EBGP CE neighbor.

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-25

In this example, a route map matching a specific SOO value was defined using the **ip extcommunity-list** command to establish a SOO filter. The **route-map** command was used to define the route map based on the filter.

The newly defined route map is then applied to a BGP neighbor (CE router) on the PE router.

Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- **Use the address-family ipv4 vrf-name command in the BGP routing process to configure a per-VRF BGP routing context.**
- **Service providers offering MPLS VPN services are at risk of denial-of-service attacks. Limiting VRF tables is one method to prevent such attacks.**
- **Use the neighbor maximum-prefix command to limit the number of prefixes received from a BGP neighbor.**
- **Use the maximum router command to limit the total number of VRF routes.**

© 2004 Cisco Systems, Inc. All rights reserved. MPLS v2.1—5-26

Summary (Cont.)

Cisco.com

- **BGP loop detection prevents customers from reusing their AS number. AS-override prevents this issue by replacing the customer AS number with the ISP AS number.**
- **A customer site cannot link two VPN sites of the same AS number because of BGP loop detection. Allowas-in disables the BGP path check and permits routing updates.**
- **The SOO extended BGP community is used as a loop prevention mechanism for multihomed customer sites.**

© 2004 Cisco Systems, Inc. All rights reserved. MPLS v2.1—5-27

Troubleshooting MPLS VPNs

Overview

This lesson explains the preliminary steps for troubleshooting an MPLS VPN. The lesson also looks at routing information flow troubleshooting and VPN data flow troubleshooting.

It is important to be able to determine what steps you should take when trying to solve a problem with your MPLS VPN network. This lesson looks at how to go about correcting MPLS VPN network problems.

Objectives

Upon completing this lesson, you will be able to describe how to troubleshoot MPLS VPN operations. This ability includes being able to meet these objectives:

- Identify the preliminary steps in MPLS VPN troubleshooting
- Identify the issues that you should consider when verifying the routing information flow in an MPLS VPN
- Describe the process used to validate CE-to-PE routing information flow
- Describe the process used to validate PE-to-PE routing information flow
- Describe the process used to validate PE-to-CE routing information flow
- Identify the issues that you should consider when verifying the data flow in an MPLS VPN
- Describe how to validate CEF status
- Describe how to validate the end-to-end label-switched path
- Describe how to validate the LFIB status

Identifying Preliminary Steps in MPLS VPN Troubleshooting

This topic identifies the preliminary steps in MPLS VPN troubleshooting.

Preliminary Steps in MPLS VPN Troubleshooting

Cisco.com

Perform basic MPLS troubleshooting:

- **Is CEF enabled?**
- **Are labels for IGP routes generated and propagated?**
- **Are large labeled packets propagated across the MPLS backbone (maximum transmission unit issues)?**

© 2004 Cisco Systems, Inc. All rights reserved.MPLS v2.1—5-3

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

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

MPLS VPN troubleshooting consists of these two major steps:

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

Verifying the Routing Information Flow

This topic identifies the issues that you should consider when verifying the routing information flow in an MPLS VPN.

Verifying the Routing Information Flow

Cisco.com

Verify the routing information flow:

- Are CE routes received by a PE router?
- Are routes redistributed into MP-BGP with proper extended communities?
- Are VPNv4 routes propagated to other PE routers?
- Is the BGP route selection process working correctly?
- Are VPNv4 routes inserted into VRFs on other PE routers?
- Are VPNv4 routes redistributed from BGP into the PE-CE routing protocol?
- Are VPNv4 routes propagated to other CE routers?

© 2004 Cisco Systems, Inc. All rights reserved.MPLS v2.1—5-4

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

Validating CE-to-PE Routing Information Flow

This topic describes the process that is used to validate CE-to-PE routing information flow.

Validating CE-to-PE Routing Information Flow

Cisco.com

Are CE routes received by the PE router?

- **Verify with the `show ip route vrf vrf-name` command on PE-1.**
- **Perform traditional routing protocol troubleshooting if needed.**

© 2004 Cisco Systems, Inc. All rights reserved. MPLS v2.1—5-5

Troubleshooting routing information flow requires the verification of end-to-end routing information propagation between CE routers. The first step is to check the routing information exchange from CE routers to PE routers. Use the **show ip route vrf vrf-name** command to verify that the PE router receives customer routes from the CE router. Use traditional routing protocol troubleshooting if needed. Troubleshooting of standard enterprise routing protocols is described in the *Cisco Internetwork Troubleshooting (CIT)* course. BGP-specific troubleshooting is described in the individual modules of the *Configuring BGP on Cisco Routers (BGP)* course.

Validating PE-to-PE Routing Information Flow

This topic describes the process that is used to validate PE-to-PE routing information flow.

Validating PE-to-PE Routing Information Flow

Cisco.com

The diagram illustrates a P-Network (Private Network) containing two PE (Provider Edge) routers, PE-1 and PE-2. PE-1 is connected to two CE-Spoke (Customer Edge) routers. PE-2 is also connected to two CE-Spoke routers. A red curved arrow on PE-1 indicates a loopback configuration. A green arrow points from one of the CE-Spoke routers connected to PE-1 towards PE-1, representing the flow of routing information. The Cisco logo is visible in the bottom right corner of the diagram area.

Are routes redistributed into MP-BGP with proper extended communities?

- **Verify with the `show ip bgp vpnv4 vrf vrf-name ip-prefix` command on PE-1.**
- **Troubleshoot with debug ip bgp commands.**

© 2004 Cisco Systems, Inc. All rights reserved. MPLS v2.1—5-6


The CE routes received by the PE router need to be redistributed into MP-BGP; otherwise, they will not get propagated to other PE routers. Common configuration mistakes in this step include the following:

- Failing to configure redistribution between the PE-CE routing protocol and the per-VRF routing context of the BGP
- Using a route map on redistribution that filters CE routes

Proper redistribution of CE routes into a per-VRF instance of BGP can be verified with the **show ip bgp vpnv4 vrf vrf-name** command. The RD prepended to the IPv4 prefix and the RTs attached to the CE route can be verified with the **show ip bgp vpnv4 vrf vrf-name ip-prefix** command.

Validating PE-to-PE Routing Information Flow (Cont.)

Cisco.com



Are VPNv4 routes propagated to other PE routers?

- Verify with the `show ip bgp vpnv4 all ip-prefix/length` command.
- Troubleshoot PE-to-PE connectivity with traditional BGP troubleshooting tools.

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-7


The CE routes redistributed into MP-BGP need to be propagated to other PE routers. Verify proper route propagation with the `show ip bgp vpnv4 all ip-prefix` command on the remote PE router.

Note Routes sent by the originating PE router might not be received by a remote PE router because of automatic RT-based filters installed on the remote PE router.

Automatic route filters are based on RTs. Verify that the RTs attached to the CE route in the originating PE router match at least one of the RTs configured as import RTs in the VRF on the receiving PE router.

Validating PE-to-PE Routing Information Flow (Cont.)

Cisco.com



Is the BGP route selection process working correctly on PE-2?

- **Verify with the `show ip bgp vpnv4 vrf vrf-name ip-prefix` command.**
- **Change local preference or weight settings if needed.**
- **Do not change MED if you are using IGP-BGP redistribution on PE-2.**

© 2004 Cisco Systems, Inc. All rights reserved.


MPLS v2.1—5-8

In complex environments with multihomed customer sites, the BGP route selection process might affect proper MPLS VPN operation. Use standard BGP route selection tools (weights or local preference) to influence BGP route selection. The MED attribute should not be changed inside the MPLS VPN backbone if you plan to use two-way route redistribution between the PE-CE routing protocol and BGP.

Refer to the BGP course for more information on BGP weights, local preference, and the MED attribute.

Validating PE-to-PE Routing Information Flow (Cont.)

Cisco.com



Are VPNv4 routes inserted into VRFs on PE-2?

- Verify with the `show ip route vrf` command.
- Troubleshoot with the `show ip bgp ip-prefix` and `show ip vrf detail` command.
- Perform additional BGP troubleshooting if needed.

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-9

The VPNv4 routes received by the PE router have to be inserted into the proper VRF. This insertion can be verified with the `show ip route vrf` command. Common configuration mistakes in this step include the following:


- The wrong import RTs are configured in the VRF.
- The route map configured as the import route map is rejecting the VPNv4 routes. Refer to the “Using Advanced VRF Import and Export Features” lesson in the “Complex MPLS VPNs” module for more information on import route maps.

The validity of the import RTs can be verified with the `show ip bgp vpnv4 all ip-prefix` command, which displays the RTs attached to a VPNv4 route. You can also verify the validity of the import RTs with the `show ip vrf detail` command, which lists the import RTs for a VRF. At least one RT attached to the VPNv4 route needs to match at least one RT in the VRF.

Note Be patient when troubleshooting this step. The import of VPNv4 routes into VRFs is not immediate and can take more than a minute in the worst circumstances.

Validating PE-to-PE Routing Information Flow (Cont.)

Cisco.com



Are VPNv4 routes redistributed from BGP into the PE-CE routing protocol?

- Verify redistribution configuration—is the IGP metric specified?
- Perform traditional routing protocol troubleshooting.

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-10

Finally, the BGP routes received via MP-BGP and inserted into the VRF need to be redistributed into the PE-CE routing protocol. A number of common redistribution mistakes can occur here, starting with missing redistribution metrics.

Refer to the *Building Scalable Cisco Internetworks (BSCI)* and *Cisco Internetwork Troubleshooting (CIT)* courses for more information on route redistribution troubleshooting.

Validating PE-to-CE Routing Information Flow

This topic describes the process used to validate PE-to-CE routing information flow.

Validating PE-to-CE Routing Information Flow

Cisco.com

Are VPNv4 routes propagated to other CE routers?

- **Verify with the show ip route command on CE spoke.**
- **Alternatively, do CE spokes have a default route toward PE-2?**
- **Perform traditional routing protocol troubleshooting if needed.**

© 2004 Cisco Systems, Inc. All rights reserved. MPLS v2.1—5-11

Last but not least, the routes redistributed into the PE-CE routing protocol have to be propagated to CE routers. You may also configure the CE routers with a default route toward the PE routers (see note). Use standard routing protocol troubleshooting techniques in this step.

Note When using a default route on the CE routers, verify that the CE routers use classless routing configured with the **ip classless** command.

Identifying the Issues When Verifying the Data Flow

This topic identifies the issues that you should consider when verifying the data flow in an MPLS VPN.

Verifying the Data Flow

Cisco.com

Verify proper data flow:

- Is CEF enabled on the ingress PE router interface?
- Is the CEF entry correct on the ingress PE router?
- Is there an end-to-end label switched path tunnel (LSP tunnel) between PE routers?
- Is the LFIB entry on the egress PE router correct?

© 2004 Cisco Systems, Inc. All rights reserved.MPLS v2.1—5-12

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

Validating CEF Status

This topic describes how to validate CEF status.

Validating CEF Status

Cisco.com

Is CEF enabled on the ingress PE router interface?

- Verify with the `show cef interface` command.
- MPLS VPN needs CEF enabled on the ingress PE router interface for proper operation.
- CEF might become disabled because of additional features deployed on the interface.

© 2004 Cisco Systems, Inc. All rights reserved. MPLS v2.1—5-13

One of the most common configuration mistakes related to data flow is the failure to enable CEF in the ingress PE router interface. The presence of CEF can be verified with the **show cef interface** command. CEF is the only switching method that can perform per-VRF lookup and thus support MPLS VPN architecture.

Assuming that CEF is enabled on the router, here are the three common reasons for CEF configuration mistakes:

- CEF is manually disabled on an interface.
- The interface is using an encapsulation method that is not supported by CEF, such as X.25 or Multilink PPP (MLP) with interleaving.
- Another feature has been configured on the interface that disables CEF (for example, IP precedence accounting).

Validating CEF Status: show cef interface

Cisco.com

```
Router#show cef interface serial 1/0.20
Serial1/0.20 is up (if_number 18)
Internet address is 150.1.31.37/30
ICMP redirects are always sent
Per packet loadbalancing is disabled
IP unicast RPF check is disabled
Inbound access list is not set
Outbound access list is not set
IP policy routing is disabled
Interface is marked as point to point interface
Hardware idb is Serial1/0
Fast switching type 5, interface type 64
IP CEF switching enabled
IP CEF VPN Fast switching turbo vector
VPN Forwarding table "SiteA2"
Input fast flags 0x1000, Output fast flags 0x0
ifindex 3(3)
Slot 1 Slot unit 0 VC -1
Transmit limit accumulator 0x0 (0x0)
IP MTU 1500
```

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-14

show cef interface

To display detailed CEF information for all interfaces, use the **show cef interface** command in EXEC mode: **show cef interface *type number* [statistics][detail]**.

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

Syntax Description

Parameter	Description
<i>type number</i>	Displays interface type and number for CEF information.
statistics	(Optional) Displays switching statistics for the line card.
detail	(Optional) Displays detailed CEF information for the specified interface type and number.

Usage Guidelines

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

The **detail** keyword displays more CEF information for the specified interface. You can use this command to show the CEF state on an individual interface.


This table describes the fields shown in the output.

Syntax Description

Parameter	Description
interface type number is {up down}	Indicates status of the interface.
Internet address	Internet address of the interface.
ICMP redirects are {always sent never sent}	Indicates how packet forwarding is configured.
Per-packet load balancing	Status of load balancing in use on the interface (enabled or disabled).
IP unicast RPF check	Indicates status of IP unicast Reverse Path Forwarding (RPF) check on the interface.
Inbound access list {# Not set}	Number of inbound access lists defined for the interface.
Outbound access list	Number of outbound access lists defined for the interface.
IP policy routing	Indicates the status of IP policy routing on the interface.
Hardware idb is type number	Interface type and number configured.
Fast switching type	Indicates switching mode in use. Used for troubleshooting.
IP CEF switching {enabled disabled}	Indicates the switching path used.
Slot n Slot unit n	Slot number.
Hardware transmit queue	Indicates the number of packets in the transmit queue.
Transmit limit accumulator	Indicates the maximum number of packets allowed in the transmit queue.
IP MTU	Value of the MTU size set on the interface.

Validating CEF Status (Cont.)

Cisco.com



Is the CEF entry correct on the ingress PE router?

- Display the CEF entry with the `show ip cef vrf vrf-name ip-prefix/length detail` command.
- Verify the label stack in the CEF entry.

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—5-15

If CEF switching is enabled on the ingress interface, you can verify the validity of the CEF entry and the associated label stack with the **show ip cef vrf vrf-name ip-prefix detail** command. The top label in the stack should correspond to the BGP next-hop label as displayed by the **show tag forwarding** command on the ingress router. The second label in the stack should correspond to the label allocated by the egress router. You can verify this by using the **show tag forwarding** command on the egress router.

Validating the End-to-End LSP

This topic describes how to validate the end-to-end label-switched path (LSP).

Validating the End-to-End Label Switched Path

Cisco.com

Is there an end-to-end LSP tunnel between PE routers?

- Check summarization issues—BGP next hop should be reachable as 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.

© 2004 Cisco Systems, Inc. All rights reserved. MPLS v2.1—5-16

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

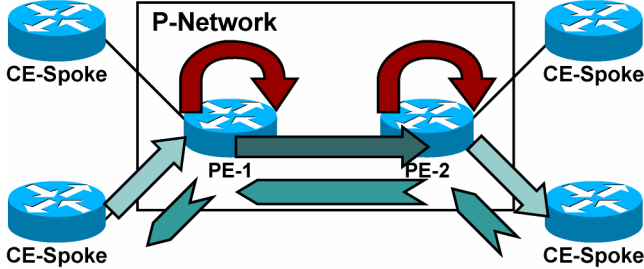
The quickest way to diagnose summarization problems is to disable IP time-to-live (TTL) propagation into the MPLS label header using the **no mpls ip ttl-propagate** command. The **tracert** command toward the BGP next hop should display no intermediate hops when TTL propagation is disabled. If intermediate hops are displayed, the LSP tunnel between PE routers is broken at those hops and the VPN traffic cannot flow.

Validating the LFIB Status

This topic describes how to validate the LFIB status.

Validating the LFIB Status

Cisco.com



Is the LFIB entry on the egress PE router correct?

- Find out the second label in the label stack on PE-2 with the `show ip cef vrf vrf-name ip-prefix detail` command.
- Verify correctness of LFIB entry on PE-1 with the `show mpls forwarding vrf vrf-name value detail` command.

© 2004 Cisco Systems, Inc. All rights reserved. MPLS v2.1—5-17

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

Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- **MPLS troubleshooting can be divided into two main steps:**
 - Verify routing information flow
 - Verify proper data flow
- **Verification of the routing information flow should be done systematically, starting at the ingress CE router and moving to the egress CE router.**
- **The first step in validating CE-to-PE routing information flow is to check the routing information exchange from CE routers to PE routers.**
- **Use the `show ip bgp vpnv4 vrf vrf-name ip-prefix` command to validate PE-to-PE routing information flow.**
- **When validating PE-to-CE routing information flow, ensure that routes are redistributed back into CE routing protocol on the PE route and propagated toward CE routers.**

© 2004 Cisco Systems, Inc. All rights reserved. MPLS v2.1—5-18

Summary (Cont.)

Cisco.com

- **Verification of the data flow should be done systematically, starting at ingress the CE router and moving to the egress CE router. Verify that CEF and LSP switching are operational.**
- **Use the `show cef interface` command to verify the CEF status.**
- **When validating the end-to-end LSP, verify that there is an end-to-end LSP tunnel between PE routers.**
- **To validate the LFIB status, verify the contents of the LFIB on the egress PE router and compare them with the second label in the label stack on the ingress PE router.**

© 2004 Cisco Systems, Inc. All rights reserved. MPLS v2.1—5-19

Module Summary

This topic summarizes the key points discussed in this module.

Module Summary

Cisco.com

- **The VRF table is a virtual routing and forwarding instance separating sites with the same connectivity requirements.**
- **Configuring VRF tables requires defining the VRF name, RD, and import and export RTs.**
- **MP-BGP configuration must define the neighbors, address family for VPNv4 routing, and finally activate the neighbors.**
- **RIPv2 and EIGRP routing for PE to CE requires use of the address-family ipv4 command to define the routing context. Redistribution is also required between IGP and BGP.**

© 2004 Cisco Systems, Inc. All rights reserved. MPLS v2.1—5-5

Module Summary (Cont.)

Cisco.com

- **Monitoring MPLS VPN operations includes monitoring MP-BGP neighbor status, VRF routing process, and CEF and LFIB status.**
- **The MPLS VPN routing model implements MP-BGP as the superbackbone for OSPF. PE to CE OSPF routing is defined as a new routing process via the VRF name.**
- **The MPLS VPN routing model implements BGP routing between PE and CE as BGP instances using the command address-family ipv4.**
- **MPLS VPN troubleshooting uses a systematic process starting at the ingress PE router and moving to the egress PE router using a variety of show commands.**

© 2004 Cisco Systems, Inc. All rights reserved. MPLS v2.1—5-5

An MPLS VPN implementation involves VRF tables, the interaction between customer-to-provider routing protocols and MP-BGP in the service provider backbone.

References

For additional information, refer to these resources:

- Cisco.com for additional information about VPNs
- “BGP Filtering and Route Selection” module in the *Configuring BGP on Cisco Routers* (BGP) course
- “Advanced BGP Configuration” module in the *Configuring BGP on Cisco Routers* (BGP) course
- *Building Scalable Cisco Networks* (BSCN) course
- *Cisco Internetwork Troubleshooting* (CIT) course

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: Using MPLS VPN Mechanisms of Cisco IOS Platforms)
- A) the routing and forwarding instance for all sites belonging to a single customer
 - B) the routing and forwarding instance for all sites belonging to a single customer location
 - C) the routing and forwarding instance for all sites using a common routing protocol
 - D) the routing and forwarding instance for a set of sites with identical connectivity requirements
- Q2) Why are VRFs used to establish separate routing protocol contexts? (Source: Using MPLS VPN Mechanisms of Cisco IOS Platforms)
-
-
- Q3) Which two protocols are VPN-aware? (Choose two.) (Source: Using MPLS VPN Mechanisms of Cisco IOS Platforms)
- A) RIPv2
 - B) IS-IS
 - C) ODR
 - D) EIGRP
- Q4) True or False? VRFs are assigned to an interface. (Source: Using MPLS VPN Mechanisms of Cisco IOS Platforms)
-
-
- Q5) A PE router is supporting site A for a VPN on one interface using RIP as the routing protocol. Site B belongs to the same VPN and is being supported on a second interface using EBGp as the routing protocol. Why is it necessary to redistribute the RIP-learned route into the per-VRF instance of the BGP process? (Source: Using MPLS VPN Mechanisms of Cisco IOS Platforms)
- A) to allow site A and B to communicate with each other
 - B) to allow the RIP route to be propagated to the VRF routing tables
 - C) to allow the RIP routes to be propagated to the local EBGp session
 - D) to allow the RIP routes to be propagated through the backbone MP-BGP process to other PE routers

- Q6) How are VPNv4 routers propagated to a RIP-speaking CE router? (Source: Using MPLS VPN Mechanisms of Cisco IOS Platforms)
-
-
-
- Q7) Which command do you use to create a VRF named VPNA? (Source: Configuring VRF Tables)
- A) **ip vrf VPNA**
 - B) **ip rt vrf VPNA**
 - C) **ip rd vrf VPNS**
 - D) **ip vrf forwarding VPNA**
- Q8) Which two VRF parameters specify the extended community attribute used in VPNv4 BGP? (Choose two.) (Source: Configuring VRF Tables)
- A) **rd route-distinguisher**
 - B) **route-target export RT**
 - C) **route-target import RT**
 - D) **ip vrf forwarding vrf-name**
- Q9) Which command do you use to associate interface e0/0 with a VRF named VPNA? (Source: Configuring VRF Tables)
- A) **ip vrf VPNA**
 - B) **ip vrf VPNA int e0/0**
 - C) **ip vrf forwarding VPNA**
 - D) **ip vrf VPNA forwarding e0/0**
- Q10) What happens to the interface of an existing IP address when you associate the interface with a VRF? (Source: Configuring VRF Tables)
- A) It will remain unchanged.
 - B) It will be removed from the interface.
 - C) It will be changed to the loopback 0 address.
 - D) It will be moved under the VRF configuration.
- Q11) You have created a configuration that defines three import route targets (650001:01, 650002:02, and 650003:03) for a VRF. A route update has three RTs (650003:03, 650004:04, and 650005:05) attached to it. How will this update be processed and why? (Source: Configuring VRF Tables)
- A) The update will be accepted by the VRF because it matches the import RD of 03.
 - B) The update will be discarded by the VRF because it does not match all of the RTs in the import list.
 - C) The update will be accepted by the VRF because it matches at least one of the RTs in the import list.
 - D) The update will be discarded by the VRF because it does not match all of the RDs in the import list.

Q12) In which two ways does the MPLS VPN architecture use the BGP routing protocol?
(Source: Configuring an MP-BGP Session Between PE Routers)

Q13) What is a BGP address family? (Source: Configuring an MP-BGP Session Between PE Routers)

Q14) What are the two types of BGP address families that can be configured on a PE router?
(Source: Configuring an MP-BGP Session Between PE Routers)

Q15) Which mandatory parameters do you have to configure on an MP-BGP neighbor?
(Source: Configuring an MP-BGP Session Between PE Routers)

Q16) Why is it necessary to enable extended BGP communities when you are supporting MPLS VPNs? (Source: Configuring an MP-BGP Session Between PE Routers)

Q17) Why would you want to disable propagation of IPv4 routing updates between MP-BGP neighbors? (Source: Configuring an MP-BGP Session Between PE Routers)

Q18) How do you configure the routing context in RIP? (Source: Configuring Small Scale Routing Protocols Between PE and CE Routers)

Q19) How do you propagate static VRF routes between PE routers? (Source: Configuring Small Scale Routing Protocols Between PE and CE Routers)

Q20) How would you configure redistribution to propagate customer RIP routing updates across the MPLS VPN backbone? (Source: Configuring Small Scale Routing Protocols Between PE and CE Routers)

Q21) Which three commands do you use to display all configured VRFs on the router? (Source: Monitoring MPLS VPN Operations)

Q22) How do you verify the contents of a VRF routing table? (Source: Monitoring MPLS VPN Operations)

Q23) Why is the BGP protocol always running in every VRF and how would you display the BGP parameter related to a VRF? (Source: Monitoring MPLS VPN Operations)

Q24) How do you verify that a session has been established between two VPNv4 neighbors?
(Source: Monitoring MPLS VPN Operations)

Q25) How do you verify the contents of a BGP VPNv4 routing table? (Source: Monitoring MPLS VPN Operations)

Q26) Which three commands can be used to display per-VRF FIB and LFIB information?
(Source: Monitoring MPLS VPN Operations)

Q27) Which command can be used to display tags assigned to local or remote VRF routes by the local or remote PE router? (Source: Monitoring MPLS VPN Operations)

Q28) Which command do you use to perform each of the following traceroutes? (Source: Monitoring MPLS VPN Operations)

Ingress CE to egress PE: _____

Ingress CE to egress CE: _____

Ingress PE to egress PE: _____

Ingress PE to egress CE: _____

Ingress P to egress PE: _____

Ingress P to egress CE: _____

Q29) Why is the OSPF superbackbone needed in MPLS VPN environments? (Source: Configuring OSPF as the Routing Protocol Between PE and CE Routers)

Q30) What is the interaction between Area 0 and a superbackbone? (Source: Configuring OSPF as the Routing Protocol Between PE and CE Routers)

Q31) What is the interaction between a superbackbone and other areas? (Source: Configuring OSPF as the Routing Protocol Between PE and CE Routers)

Q32) How are OSPF route attributes propagated across an MPLS VPN backbone? (Source: Configuring OSPF as the Routing Protocol Between PE and CE Routers)

Q33) What is the purpose of the down bit in an LSA header? (Source: Configuring OSPF as the Routing Protocol Between PE and CE Routers)

Q34) Why do you need a VRF route limit command? (Source: Configuring BGP as the Routing Protocol Between PE and CE Routers)

Q35) When would you need the AS-override feature? (Source: Configuring BGP as the Routing Protocol Between PE and CE Routers)

Q36) How does the AS-override feature work? (Source: Configuring BGP as the Routing Protocol Between PE and CE Routers)

Q37) When would you need the allowas-in feature? (Source: Configuring BGP as the Routing Protocol Between PE and CE Routers)

Q38) When is it necessary to use the AS-override feature instead of the allowas-in feature? (Source: Configuring BGP as the Routing Protocol Between PE and CE Routers)

Q39) How do you prevent BGP loops when using AS-override? (Source: Configuring BGP as the Routing Protocol Between PE and CE Routers)

Q40) How do you prevent BGP loops when using allowas-in? (Source: Configuring BGP as the Routing Protocol Between PE and CE Routers)

Q41) What is the Site of Origin? (Source: Configuring BGP as the Routing Protocol Between PE and CE Routers)

Q42) When would you have to use the Site of Origin? (Source: Configuring BGP as the Routing Protocol Between PE and CE Routers)

Q43) Where can you set the Site of Origin? (Source: Configuring BGP as the Routing Protocol Between PE and CE Routers)

Q44) What are the preliminary MPLS VPN troubleshooting steps? (Source: Troubleshooting MPLS VPNs)

Q45) Which command do you use to verify that the PE router is receiving customer routes from the CE router? (Source: Troubleshooting MPLS VPNs)

Q46) How do you verify the routing information exchange between PE routers? (Source: Troubleshooting MPLS VPNs)

Q47) How do you verify redistribution of VPNv4 routes into the PE-CE routing protocol? (Source: Troubleshooting MPLS VPNs)

Q48) How do you test end-to-end data flow between PE routers? (Source: Troubleshooting MPLS VPNs)

Q49) How do you verify that the PE router ingress interface supports CEF switching? (Source: Troubleshooting MPLS VPNs)

Q50) How do you verify that there is an end-to-end LSP? (Source: Troubleshooting MPLS VPNs)

Q51) How do you verify that the LFIB entry on the egress PE router is correct? (Source: Troubleshooting MPLS VPNs)

Module Self-Check Answer Key

- Q1) D
- Q2) Because with routing protocols such as RIP and BGP, only a single copy of the protocol may be running in the router.
- Q3) A, D
- Q4) False. Interfaces are assigned to a VRF.
- Q5) D
- Q6) VPNv4 routes are redistributed from the global BGP table to the per-instance BGP table and then to the per-instance RIP, which is propagated to the CE router.
- Q7) A
- Q8) B, C
- Q9) C
- Q10) B
- Q11) C
- Q12) EBGP is used to carry routing updates between the PE router and the CE router.
IBGP (VNPv4) is used to carry VPN route updates between PE routers.

Note IBGP (IPv4) is used to carry *non-VPN* route updates between PE routers.

- Q13) A BGP address family is a routing protocol context that is used to configure global BGP routing, VPN routing, and CE-to-PE routing into the same BGP process.
- Q14) VPNv4 and IPv4
- Q15) **neighbor ip-address remote-as as-number**
neighbor ip-address update-source interface-type interface number
address-family vpnv4
neighbor ip-address activate
neighbor ip-address next-hop-self
- Q16) Extended BGP communities attached to VPNv4 prefixes have to be exchanged between MP-BGP neighbors because they contain the RT information.
- Q17) when you are supporting only VPN routes
- Q18) On the CE, enable RIP. On the PE, enable RIP and use the **address-family ipv4 vrf vrf-name** command under the router RIP section.
- Q19) by enabling the static route using the **ip route vrf name static route parameters** command
- Q20) By using the **redistribute rip** command under the BGP address family on the ingress PE router to redistribute the RIP updates into MP-BGP. MP-BGP uses VPNv4 updates to propagate the updates to the egress PE router. The **redistribute bgp metric transparent** command under the RIP address family is used on the egress PE to redistribute the updates back into RIP.
- Q21) **show ip vrf**
show ip vrf detail
show ip vrf interfaces
- Q22) Use the **show ip route vrf** command.

- Q23) The BGP protocol is needed to carry the VPNv4 routes. Use the **show ip bgp vpnv4 vrf** command.
- Q24) Use the **show ip bgp neighbors** command and verify that the VPNv4 status is “advertised and received.”
- Q25) Use the **show ip bgp vpnv4 vrf vrf-name** command.
- Q26) **show ip cef vrf**
show ip cef vrf detail
show mpls forwarding vrf
- Q27) **show ip bgp vpnv4 tags**
- Q28) Ingress CE to egress PE: **trace**
Ingress CE to egress CE: **trace**
Ingress PE to egress PE: **trace**
Ingress PE to egress CE: **trace vrf vrf-name**
Ingress P to egress PE: **trace**
Ingress P to egress CE: You cannot do a traceroute from a P router to any CE router. The P router does not have the CE routing information in its routing table.
- Q29) Because MPLS VPNs use BGP to propagate routes between sites; internal OSPF routers in one area will appear as external routes in another area unless the superbackbone makes the MPLS VPN backbone transparent to OSPF.
- Q30) The superbackbone is transparent to Area 0.
- Q31) The superbackbone appears as Area 0 to non-Area 0 areas.
- Q32) The OSPF routes are propagated into BGP.
The OSPF metrics and LSA information are carried in the BGP community attribute.
- Q33) The down bit is used to prevent routing loops.
- Q34) You need a VRF route limit command because of tight coupling of the customer and the service provider network in the MPLS VPN architecture. This tight coupling might also result in the service provider network being exposed to design and configuration errors in customer networks and to a number of new denial-of-service attacks based on routing protocol behavior.
- Q35) when you need to connect two or more sites that use the same AS number via a VPN
- Q36) All leading occurrences of the AS number of the receiving BGP router are replaced with the AS number of the sending BGP router. Any other occurrences (farther down the AS path) of the AS number of the receiving router are not replaced because they indicate a real routing information loop.
- Q37) In some security-conscious implementations, customer VPNs are linked by a customer router that performs security functions, such as access filtering or access logging.
- Q38) in solutions where customer VPNs are linked by a customer router that do not support the AS-override feature
- Q39) Only the leading occurrences of the AS number of the receiving BGP router are replaced with the AS number of the sending BGP router. Any other occurrences (farther down the AS path) of the AS number of the receiving router are not replaced because they indicate a real routing information loop.
- Q40) Allowas-in specifies the number of times to allow advertisement of an AS number of a PE router. Valid values are from 1 to 10 times using the *number* parameter of the **allowas-in** command.
- Q41) The SOO is an extended BGP community that is used to indicate the site that has originated the routing update.
- Q42) The SOO is used as an additional loop prevention mechanism in scenarios when the allowas-in feature is enabled.

- Q43) For routes received from BGP-speaking CE routers, the SOO is configured by the incoming route map on the PE router. For all other routes, a route map setting the SOO is applied to the incoming interface and the SOO is attached to the BGP route when an IGP route received through that interface is redistributed into BGP.
- Q44) Is CEF enabled?
Are labels for IGP routes generated and propagated?
Are large labeled packets propagated across the MPLS backbone (MTU issues)?
- Q45) **show ip route vrf** *vrf-name*
- Q46) Use the **show ip bgp vpnv4 all** *ip-prefix/length* command to verify proper route propagation.
- Q47) Use the **show ip bgp vrf** *vrf-name ip-prefix* command on the egress PE router or use the **show ip route** command on the egress CE router.
- Q48) From the ingress PE router, use the **ping vrf** *vrf-name* command to ping the interface that supports the CE router.
- Q49) Use the **show cef interface** command.
- Q50) Check LFIB values hop by hop or use the **trace vrf** *vrf-name* command from the ingress PE router.
- Q51) Find out the second label in the label stack on the ingress PE with the **show ip cef vrf** *vrf-name ip-prefix detail* command. Verify the correctness of the LFIB entry on the egress PE with the **show mpls forwarding vrf** *vrf-name value detail* command.

Complex MPLS VPNs

Overview

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

Module Objectives

Upon completing this module, you will be able to describe how the MPLS VPN model can be used to implement managed services and Internet access. This ability includes being able to meet these objectives:

- Configure advanced VRF import and export features
- Identify the characteristics of overlapping VPNs
- Identify the characteristics of the central services VPN solutions
- Identify the characteristics of the managed CE router service
- Describe MPLS VPN managed services

Using Advanced VRF Import and Export Features

Overview

Some virtual routing and forwarding (VRF) features allow you to be more selective with routes, by specifying which routes will or will not be added. You may also limit the number of routes that a customer can insert into the VRF instance. This lesson presents the command syntax that is used to limit each type of route and shows configuration examples of these commands.

It is important to understand how to fine-tune the MPLS VPN parameters that will enhance operation of the network. Customer service level agreements (SLAs) should be adhered to so that they provide the best possible solutions for each specific customer. This lesson looks at some key areas regarding the use of VRF import and export features.

Objectives

Upon completing this lesson, you will be able to describe how to configure advanced VRF import and export features. This ability includes being able to meet these objectives:

- Identify advanced VRF features
- Describe how to configure selective VRF imports
- Describe how to configure selective VRF exports

What Are Advanced VRF Features?

This topic identifies the advanced VRF features.

Advanced VRF Features

Cisco.com

Selective import:

- This feature allows you to specify additional criteria for importing routes into the VRF.

Selective export:

- This feature allows you to specify additional RTs attached to exported routes.

VRF route limit:

- This feature allows you to specify the maximum number of routes in a VRF to prevent memory exhaustion on PE routers or denial-of-service attacks.

© 2004 Cisco Systems, Inc. All rights reserved.MPLS v2.1—6-3

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 route target (RT) of the VRF.
- The selective export feature allows you to attach specific RTs to a subset of routes exported from a VRF. By default, the same RTs get attached to all exported routes.
- The VRF route limit feature allows you to limit the number of routes that the customer— or other provider edge (PE) routers—can insert in the VRF. This feature prevents undesirable consequences such as configuration errors or denial-of-service attacks.

Configuring Selective VRF Import

This topic describes how to configure selective VRF import.

Configuring Selective VRF Import

Cisco.com

- **VRF import criteria might be more specific than just the match on the RT—for example:**
 - **Import only routes with specific BGP attributes (community, and so on).**
 - **Import routes with specific prefixes or subnet masks (only loopback addresses).**
- **A route map can be configured in a VRF to make the route import more specific.**

© 2004 Cisco Systems, Inc. All rights reserved.MPLS v2.1—6-4

Selective route import into a VRF allows you to narrow the route import criteria. Selective route import uses a route map 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 map to match any BGP attribute of a route. These attributes include communities, local preference, multi-exit discriminator (MED), autonomous system (AS) path, and so on.

The import route map filter is combined with the RT import filter. A route has to pass the RT import filter first and then the import route map. The necessary conditions for a route to be imported into a VRF are as follows:

- At least one of the RTs attached to the route matches one of the import RTs configured in the VRF.
- The route is permitted by the import route map.

Configuring Selective VRF Import (Cont.)

Cisco.com

```
Router(config-vrf)#
```

```
import map route-map
```

- This command attaches a route map to the VRF import process.
- A route is imported into the VRF only if at least one RT attached to the route matches one RT configured in the VRF **and** the route is accepted by the route map.

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—6-5

import map

To configure an import route map for a VRF, use the **import map** command in VRF configuration submode: **import map** *route-map*.

This table describes the parameters for the **import map** command.

Syntax Description

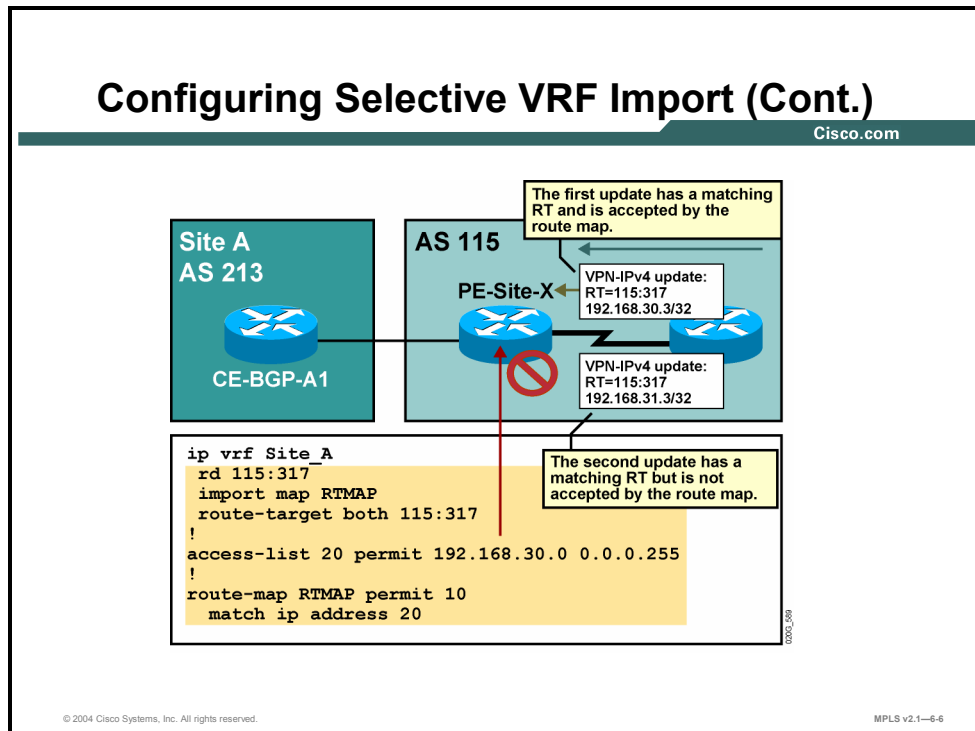
Parameter	Description
<i>route-map</i>	Specifies the route map to be used as an import route map for the VRF.

Defaults

There is no default. A VRF has no import route map unless one is configured using the **import map** command.

Example: Configuring Selective VRF Import

The figure shows an example in which an import route map is used to match the IP version 4 (IPv4) portion of incoming VPN IPv4 (VPNv4) routes and import into the VRF only routes matching a certain prefix.



A configuration similar to this one could be used to accomplish the following:

- Deploy advanced MPLS VPN topologies (for example, a managed router services topology)
- 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.

Configuring Selective VRF Export

This topic describes how to configure selective VRF export.

Configuring Selective VRF Export

Cisco.com

Routes from a VRF might have to be exported with different RTs:

- **An example would be export management routes with particular RTs.**

An export route map can be configured on VRF:

- **This route map can set extended community RTs.**
- **No other set operations can be performed by this route map.**

© 2004 Cisco Systems, Inc. All rights reserved.MPLS v2.1—6-7

Some advanced MPLS VPN topologies are easiest to implement if you can attach a variety of RTs to routes exported from the same VRF. This capability allows only a subset of the routes exported from a VRF to be imported into another VRF. Most services in which customer routers need to connect to a common server (for example, network management stations, voice gateways, and application servers) fall into this category.

The export route map function provides exactly this functionality. A route map can be specified for each VRF to attach additional RTs to routes exported from that VRF. The export route map performs only the attachment of RTs. It does not perform any filtering function.

Attributes attached to a route with an export route map are combined with the export RT attributes. If you specify export RTs in a VRF and set RTs with an export route map, all specified RTs will be attached to the exported route.

Note The export route map provides functionality almost identical to that of the import route map, but applied to a different VRF. Any requirement that can be implemented with an export route map can also be implemented with an import route map. However, the implementation of export maps can be more complicated and harder to manage.

Configuring Selective VRF Export (Cont.)

Cisco.com

Router(config)#

```
route-map name permit seq
match condition
set extcommunity rt extended-community-value [additive]
```

- This command creates a route map that matches routes based on any route map conditions and sets RTs.

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—6-8

set extcommunity

To set the BGP extended communities attribute, use the **set extcommunity** command in route-map configuration mode. To delete the entry, use the **no** form of this command.

- **set extcommunity** {rt *extended-community-value* [additive] | soo *extended-community-value*}
- **no set extcommunity** {rt *extended-community-value* [additive] | soo *extended-community-value*}

This table describes the parameters for the **set extcommunity** command.

Syntax Description

Parameter	Description
rt	Specifies the RT extended community attribute.
soo	Specifies the SOO extended community attribute.
<i>extended-community-value</i>	Specifies the value to be set. The value can be one of the following combinations: <ul style="list-style-type: none">■ autonomous-system-number: network-number■ ip-address: network-number The colon is used to separate the AS number from the network number or the IP address from the network number.
additive	(Optional) Adds an RT to the existing RT list without replacing any RTs.

Defaults

No BGP extended community attributes are set by the route map.

Configuring Selective VRF Export (Cont.)

Cisco.com

```
router(config-vrf)#  
export map name
```

- **This command attaches a route map to the VRF export process.**
- **All exported routes always get RTs configured with the route-target export command in the VRF.**
- **A route that is matched by the export route map will have additional RTs attached.**

© 2004 Cisco Systems, Inc. All rights reserved.MPLS v2.1—6-9

export map

To apply a route map to filter and modify exported routes, use the **export map** command in VRF configuration mode. To remove the route map from the VRF, use the **no** form of this command.

- **export map** *route-map*
- **no export map** *route-map*

This table describes the parameters for the **export map** command.

Syntax Description


Parameter	Description
<i>route-map</i>	Specifies the name of the route map to be used.

Defaults

No route map is used.

Example: Configuring Selective VRF Export

In the figure, the configuration is implemented with an export route map.



In the earlier example, selective import of routes into a VRF was achieved with an import route map in the receiving VRF that allowed only routes from a certain address block to be inserted into the VRF. In this example, routes from a certain address block are marked with an additional RT in the originating VRF and are automatically inserted into the receiving VRF on the basis of their RT.

The main difference between import and export route maps is therefore the deployment point, as described here:

- The import route map is deployed in the receiving VRF.
- The export route map is deployed in the originating VRF.
- Based on the network design, one or the other functionality might be preferred.

Note Import and export route maps can increase the number of routes processed by a router. The BGP maximum-prefix function can be used to ensure that the number of routes does not exceed the network design. (See the “Configuring BGP as the Routing Protocol Between PE and CE Routers” lesson in the “MPLS VPN Implementation” module for further details.)

Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- **There are three advanced VRF features: selective import, selective export, and VRF route limit.**
- **Use the import map command to configure an import route map for VRF.**
- **Use the export map command to attach a route map to the VRF export process.**

© 2004 Cisco Systems, Inc. All rights reserved. MPLS v2.1—6-11

Introducing Overlapping VPNs

Overview

Overlapping VPNs are usually used to connect parts of two separate VPNs. A third VPN is created within the MPLS VPN network that contains sites from both VPNs. A new RT extended community is used for networks originating in the sites that are also in the new VPN. This action may require a new VRF, resulting in a new route distinguisher (RD). Networks originating in these sites are exported with two RT extended communities: one for the original VPN and one for the overlapping VPN. This lesson looks at the requirements, usage, and solutions associated with overlapping VPNs.

It is important to understand customer needs and how to best fit those needs. This lesson looks at an area that helps to clarify some solutions regarding multiple separate VPNs.


Objectives

Upon completing this lesson, you will be able to identify the characteristics of overlapping VPNs. This ability includes being able to meet these objectives:

- Identify the participants in overlapping VPNs
- Identify typical overlapping VPN usages
- Describe the routing update flow in an overlapping VPN
- Describe the data flow in an overlapping VPN
- Describe how to configure overlapping VPNs

Who Are the Participants in Overlapping VPNs?

This topic identifies the participants in overlapping VPNs.



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

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

What Are Typical Overlapping VPN Usages?

This topic identifies typical overlapping VPN usages.

Typical Overlapping VPN Usages

Cisco.com

- **Companies where central sites participate in a corporate network and in an extranet**
- **A company with several security-conscious departments that exchange data between their servers**

© 2004 Cisco Systems, Inc. All rights reserved.MPLS v2.1—6-4

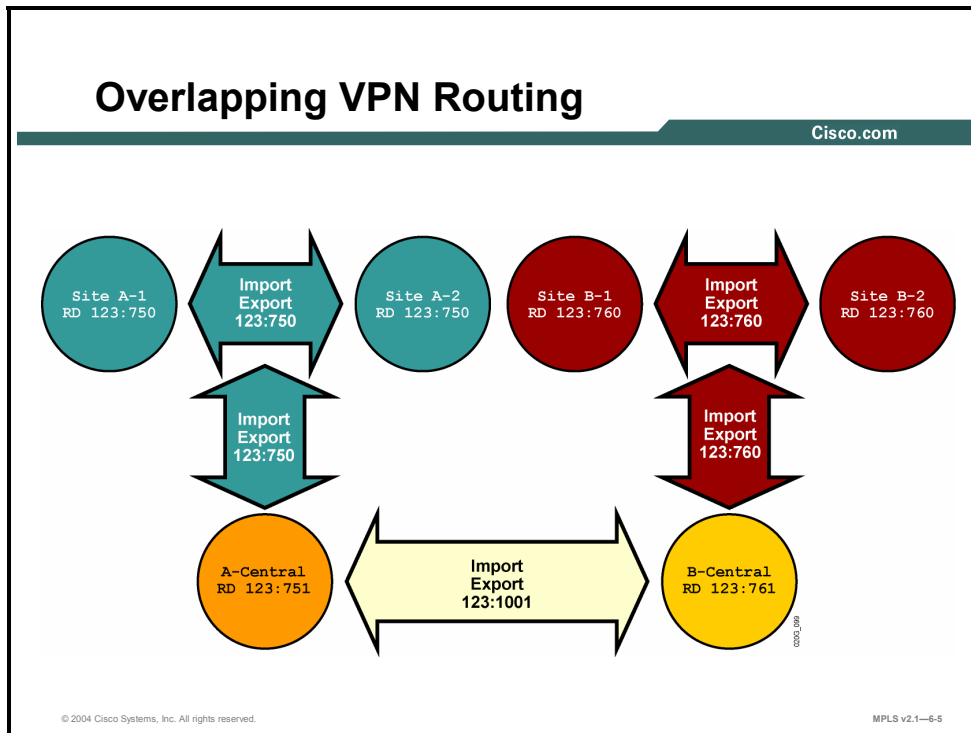
The two typical uses for overlapping VPNs are as follows:

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

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

Overlapping VPN Routing

This topic describes the routing update flow in an overlapping VPN.



Some key points regarding the routing update flow in overlapping VPNs are as follows:

- Each VPN has its own RT (123:750, 123:760) that the sites participating in the VPN import and export.
- Sites that participate in more than one VPN import routes with RTs from any VPN in which they participate and export routes with RTs for all VPNs in which they participate.

Example: Overlapping VPN Routing

The figure shows how to implement overlapping VPNs.

For site A-1 and site A-2 (participating only in VPN-A), do the following:

- Export all networks with RT 123:750
- Import all networks that carry RT 123:750 (VPN-A)

For site B-1 and site B-2 (participating only in VPN-B), do the following:

- Export all networks with RT 123:760
- Import all networks that carry RT 123:760 (VPN-B)

For site A-Central (participating in VPN-A and the overlapping VPN), do the following:

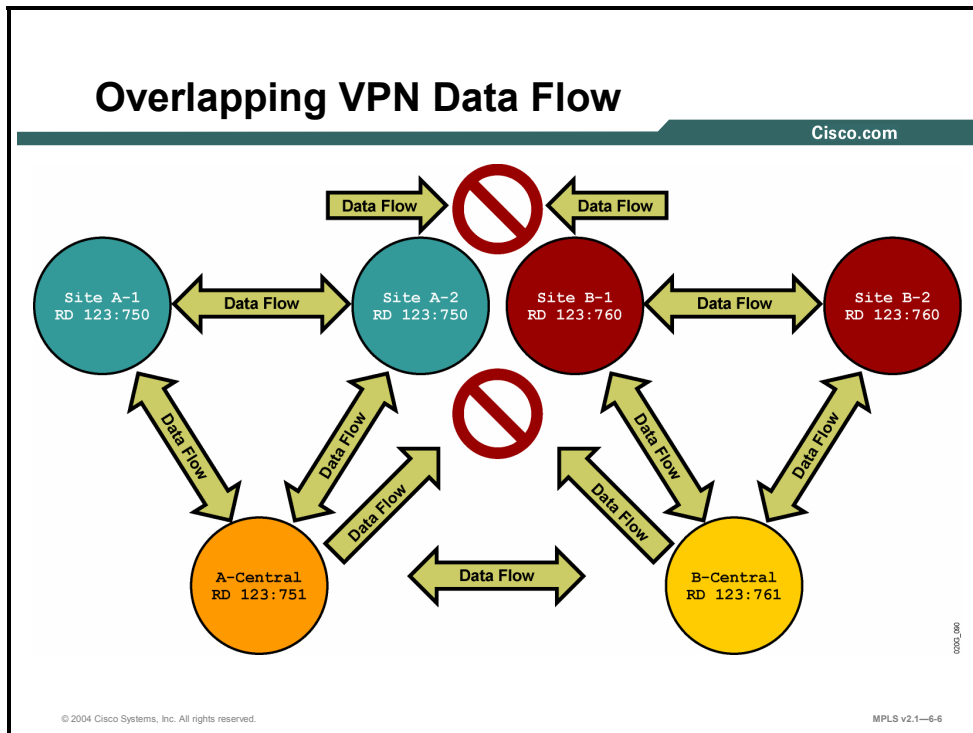
- Exports all networks with RTs 123:750 *and* 123:1001
- Imports all networks that carry RT 123:750 (VPN-A) *or* 123:1001 (overlapping VPN)

For site B-Central (participating in VPN-B and the overlapping VPN), do the following:

- Exports all networks with RTs 123:760 *and* 123:1001
- Imports all networks that carry RT 123:760 (VPN-B) *or* 123:1001 (overlapping VPN)

Overlapping VPN Data Flow

This topic describes the data flow in an overlapping VPN.




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

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

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

Configuring Overlapping VPNs

This topic describes how to configure overlapping VPNs.



You can have a network with four types of sites with different VPN memberships.

Example: Overlapping VPNs—Configuration Tasks

The figure illustrates this example. This situation requires at least the following four VRFs:

- A-Spoke-1 and A-Spoke-2 are members of VPN-A only. (They need two VRFs because they are not connected to the same PE router; they can, however, use the same RD.)
- B-Spoke-1 and B-Spoke-2 are members of VPN-B only. (They need two VRFs because they are not connected to the same PE router; they can, however, use the same RD.)
- A-Central is a member of VPN-A and overlapping VPN-AB. (They need an additional RD.)
- B-Central is a member of VPN-B and overlapping VPN-AB. (They cannot use the same RD as A-Central because B-Central has different routing requirements from A-Central.)

This table shows an RT and RD numbering scheme for PE-1.

PE-1 RT and RD Numbering Scheme

VRF	RD	Import RT	Export RT
VPN-A	123:750	123:750	123:750
VPN-B	123:760	123:760	123:760
VPN-A-Central	123:751	123:750 123:1001	123:750 123:1001


This table shows an RT and RD numbering scheme for PE-2.

PE-2 RT and RD Numbering Scheme

VRF	RD	Import RT	Export RT
VPN-A	123:750	123:750	123:750
VPN-B	123:760	123:760	123:760
VPN-B-Central	123:761	123:760 123:1001	123:760 123:1001

Configuring Overlapping VPN VRFs

Cisco.com



The Cisco IOS software configuration for PE-1 and PE-2 reflects the RT and RD numbering scheme from the two tables.

Example: Configuring Overlapping VPN VRFs

The figure shows only VRF configuration and does not show VPN routing or Multiprotocol Border Gateway Protocol (MP-BGP) routing between the provider edge (PE) routers.

Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- **Overlapping VPNs are used to provide connectivity between segments of two VPNs.**
- **There are two uses for overlapping VPNs:**
 - **Companies that use MPLS VPNs to implement both intranet and extranet services**
 - **Companies that might decide to limit visibility between departments**
- **Sites that participate in more than one (overlapping) VPN import and export routes with RTs from any VPN in which they participate.**
- **Sites cannot talk to each other if they belong to different VPNs.**
- **Overlapping VPN sites are configured with one VRF of the same RD per set of sites with the same VPN membership. RTs are configured based on the VPN membership of each respective site.**

© 2004 Cisco Systems, Inc. All rights reserved.MPLS v2.1—6-9

Introducing Central Services VPNs

Overview

A central services VPN is used when more VPNs need to share a common set of servers. These servers reside in the central services VPN, and all other VPNs have access to this VPN. Those VPNs, however, are not able to see one another. The central services VPN is implemented using two RT extended communities, where one imports networks into the VPN and the other exports networks. The client sites do the opposite. Two RT extended communities are needed to prevent client sites from exchanging routing information. This lesson looks at central services VPN solution topologies and how routing updates within that topology would flow. The lesson also discusses the implications of combining a central services VPN with a simple customer VPN. It is important to understand fully the topologies that make the most sense for the customer, and to be able to configure or recommend alternative options.

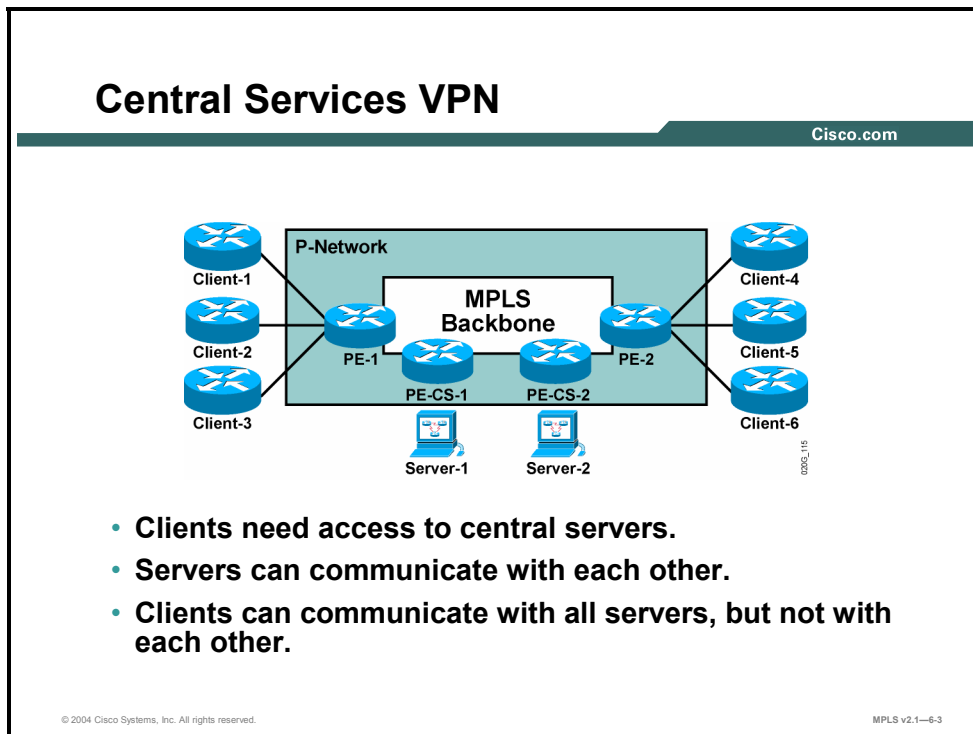
Objectives

Upon completing this lesson, you will be able to identify the characteristics of the central services VPN. This ability includes being able to meet these objectives:

- Describe the access characteristics of a central services VPN
- Describe the routing characteristics of a central services VPN
- Describe the data flow within a central services VPN
- Describe how to configure a central services VPN
- Identify the connectivity requirements when you are integrating a central services VPN with a simple VPN
- Identify the RD requirements when you are integrating a central services VPN with a simple VPN
- Identify the RT requirements when you are integrating a central services VPN with a simple VPN

What Are the Access Characteristics of a Central Services VPN?

This topic describes the access characteristics of a central services VPN.



A central services VPN is a topology with the following characteristics:

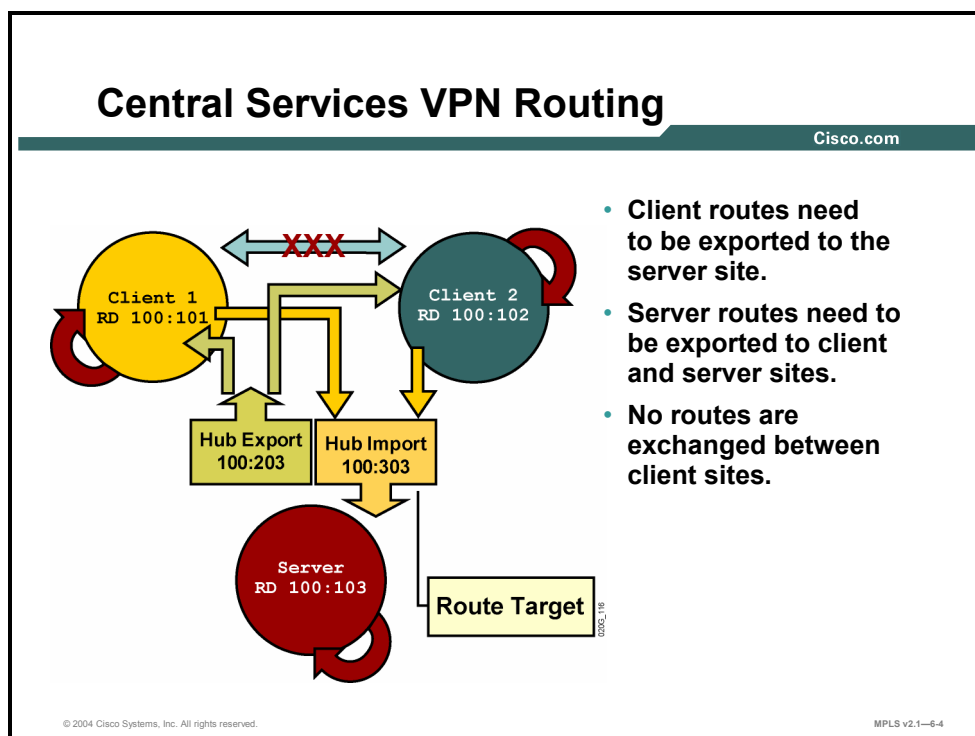
- Some sites (“server sites”) can communicate with all other sites.
- All the other sites (“client sites”) can communicate only with the server sites.

This topology can be used in the following situations:

- The service provider offers services to all customers by allowing them access to a common VPN.
- Two (or more) companies want to exchange information by sharing a common set of servers.
- A security-conscious company separates its departments and allows them access only to common servers.

What Are the Routing Characteristics of a Central Services VPN?

This topic describes the routing characteristics of a central services VPN.



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

Example: Central Services VPN Routing

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

- Client 1 and client 2 have their own RTs (100:101, 100:102) that they import and export; they also export networks with RT 100:303 and import networks with RT 100:203.

Note Client-specific RTs were introduced to comply with the implementation requirements of Cisco IOS Release 12.0 T, in which each VRF has to have at least one of its export RTs configured as its import RT.

- The central site imports and exports networks with the RT of its VPN, but it also imports networks with RT 100:303 and exports networks with RT 100:203.

Client 1 does the following:

- Export all networks with RTs 100:101 *and* 100:303
- Import all networks that carry RT 100:101 *or* 100:203

Client 2 does the following:

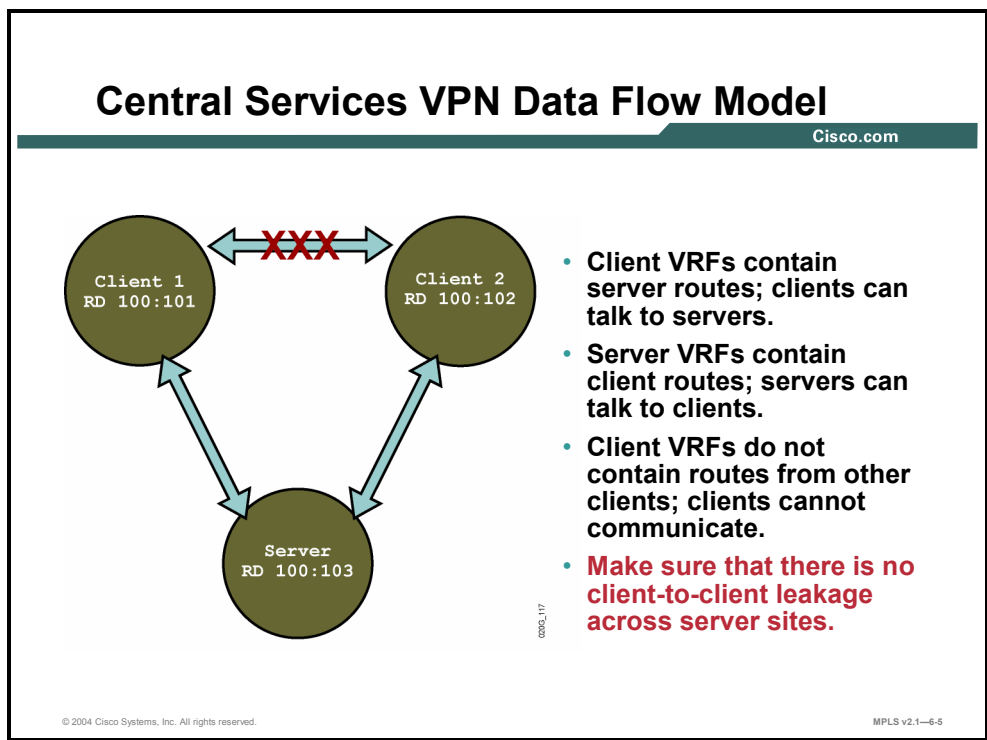
- Export all networks with RTs 100:102 *and* 100:303
- Import all networks that carry RT 100:102 *or* 100:203

The central site does the following:

- Export all networks with RT 100:203
- Import all networks that carry RT 100:303

Identifying the Central Services VPN Data Flow Model

This topic describes the data flow within a central services VPN.




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

A server VRF in this topology contains routes from the site or sites attached to the VRF, 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.

Configuring a Central Services VPN

This topic describes how to configure a central services VPN.




To configure a central services VPN, you need to address the following requirements:

- You need a separate VRF for each client.
- You need one VRF per PE router connecting a server site.
- You need a unique RD on each client site.
- You need a unique RD on each set of server sites.
- You need an import-export RT with the same value as the RD, for each client site.

Steps for Configuring a Central Services VPN (Cont.)

Cisco.com



Server sites:

- One VRF for each different service type
- Unique RD on each different service type
- Import and export routes with an RT that is the same value as the RD for each server site (VPN of server)
- Export server site routes with an RT (server-to-client)
- Import routes with the RT (clients-to-server) into the server VRFs

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—6-7

This table shows an RD and RT numbering scheme for PE-1.

PE-1 RD and RT Numbering Scheme

VRF	RD	Import RT	Export RT
Client-1	123:101	123:101	123:101
		123:203	123:303
Client-2	123:102	123:102	123:102
		123:203	123:303

This table shows an RD and RT numbering scheme for PE-2.

PE-2 RD and RT Numbering Scheme

VRF	RD	Import RT	Export RT
Client-4	123:111	123:111	123:111
		123:203	123:303
Client-5	123:112	123:112	123:112
		123:203	123:303

This table shows an RD and RT numbering scheme for PE-CS-1.


PE-CS-1 RD and RT Numbering Scheme

VRF	RD	Import RT	Export RT
Server	123:103	123:103	123:103
		123:303	123:203

This table shows an RD and RT numbering scheme for PE-CS-2.

PE-CS-2 RD and RT Numbering Scheme

VRF	RD	Import RT	Export RT
Server	123:103	123:103	123:103
		123:303	123:203



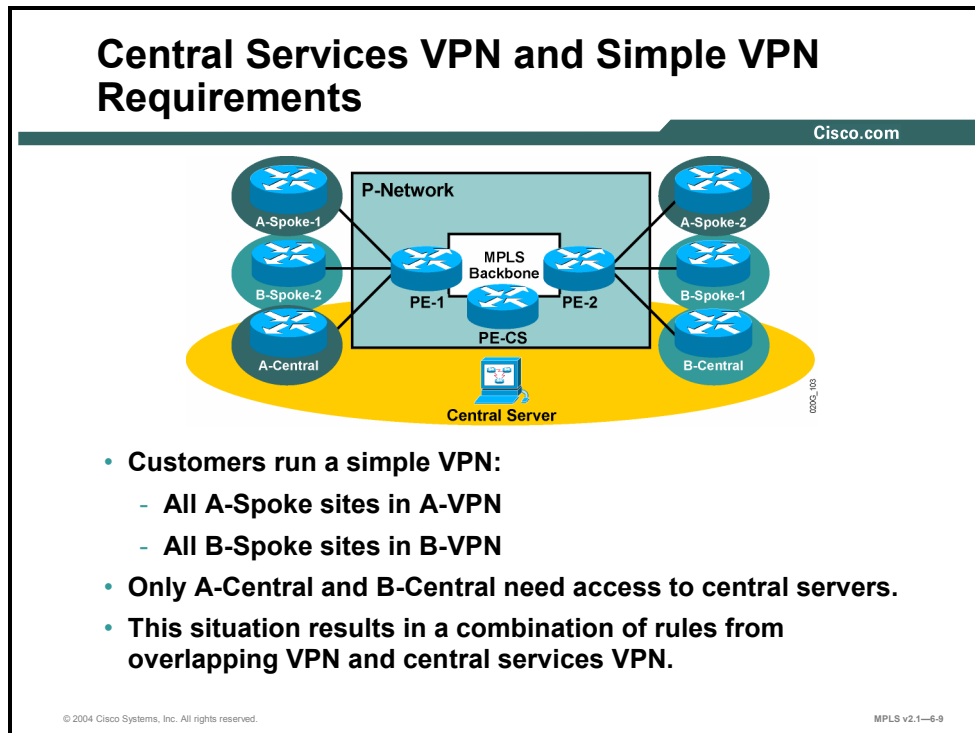
Use the **ip vrf** command to configure a central services VPN.

Example: Configuring a Central Services VPN

The figure shows a fraction of the configuration according to the RD and RT numbering scheme presented in the tables.

Integrating a Central Services VPN with a Simple VPN


This topic identifies the connectivity requirements when you are integrating a central services VPN with a simple VPN.



In this design, some of the customer sites need access to the central server. All other sites just need optimal intra-VPN access. The design is consequently a mixture of simple VPN topology and central services VPN topology.

Central Services VPN and Simple VPN Requirements (Cont.)

Cisco.com



- For all sites participating in a simple VPN, configure a separate VRF per set of sites participating in the same VPNs per PE router.
- For sites that are only clients of central servers, create a VRF per site.
- Create one VRF for central servers per PE router.

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—6-10

When integrating a central services VPN with a simple VPN, you need one VRF per VPN for sites that have access to other sites in the customer VPN but no access to the central services VPN. You need one VRF per VPN for sites that have access to the central services VPN. Finally, you need one VRF for the central services VPN; this VPN is on another PE router in the example.

Identifying the RD Requirements When Integrating a Central Services and Simple VPN

This topic identifies the RD requirements when you are integrating a central services VPN with a simple VPN.

Configuring RDs in a Central Services and Simple VPN

Cisco.com

- **Configure a unique RD for every set of VRFs with unique membership requirements:**
 - A-Spoke-1 and A-Spoke-2 can share the same RD.
 - B-Spoke-1 and B-Spoke-2 can share the same RD.
 - A-Central needs a unique RD.
 - B-Central needs a unique RD.
- **Configure one RD for all central server VRFs.**

© 2004 Cisco Systems, Inc. All rights reserved. MPLS v2.1—6-11

For this design, you need two RDs per VPN, as follows:

- One RD for simple VPN sites (The same value should also be used for import and export RTs.)
- One RD for the central services VRFs

Identifying the RT Requirements When Integrating Central Services and Simple VPN

This topic identifies the RT requirements when you are integrating a central services VPN with a simple VPN.

Configuring RTs in a Central Services and Simple VPN

Cisco.com

- **Configure the customer VPN import-export route target in all VRFs participating in customer VPN.**
- **Configure a unique import-export route target in every VRF that is only a client of central servers.**
- **Configure the central services import and export route targets in VRFs that participate in central services VPN.**

© 2004 Cisco Systems, Inc. All rights reserved. MPLS v2.1—6-12

This table shows an RD and RT numbering scheme for PE-1.

PE-1 RD and RT Numbering Scheme

VRF	RD	Import RT	Export RT
VPN-A	123:750	123:750	123:750
VPN-B	123:760	123:760	123:760
VPN-A-Central	123:751	123:750 123:101	123:750 123:100

This table shows an RD and RT numbering scheme for PE-2.

PE-2 RD and RT Numbering Scheme

VRF	RD	Import RT	Export RT
VPN-A	123:750	123:750	123:750
VPN-B	123:760	123:760	123:760
VPN-B-Central	123:761	123:760 123:101	123:760 123:100


This table shows an RD and RT numbering scheme for PE-CS.

PE-CS RD and RT Numbering Scheme

VRF	RD	Import RT	Export RT
Server	123:101	123:101 123:100	123:101

Configuring VRFs in a Central Services and Simple VPN

Cisco.com



```
ip vrf VPN A
rd 123:750
route-target both 123:750
!
ip vrf VPN A_Central
rd 123:751
route-target both 123:750
route-target export 123:100
route-target import 123:101
!
ip vrf VPN B
rd 123:760
route-target both 123:760
```

```
ip vrf Server
rd 123:101
route-target both 123:101
route-target import 123:100
```

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—6-13

Use the **ip vrf** command to configure VRFs in a central services and simple VPN.

Example: Configuring VRFs in a Central Services and Simple VPN

The example shows a fraction of the configuration according to the RD and RT numbering scheme presented in the tables.

Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- **A central services VPN is used to provide access from centralized servers to one or more customers.**
- **A central services VPN routing model indicates that:**
 - Client routes need to be exported to the server site.
 - Service routes need to be exported to client and server sites.
 - No routes are exchanged between client sites.
- **The data flow in a central services VPN model indicates that:**
 - Client VRFs contain server routes and do not contain routes from other clients.
 - Server VRFs contain client routes.
- **Some of the requirements to configure a central services VPN are:**
 - Separate VRF for each client
 - Unique RD on each client site
 - Unique RD in each set of server sites
 - Import and export RT matching between server and client sites

© 2004 Cisco Systems, Inc. All rights reserved. MPLS v2.1—6-14

Summary (Cont.)

Cisco.com

- **The hybrid of a simple VPN and a central VPN provides customers with intra-VPN access including their central site. The central sites of each customer can access centralized servers available to multiple customers.**
- **Intra-VPN customer sites can share the same RD; however, the central site of a customer and shared centralized servers require a unique RD.**
- **The import-export RT must match from respective customer intra-VPN sites to a central site. A different import-export RT set must match from the central site of the respective customers to the shared centralized server site.**

© 2004 Cisco Systems, Inc. All rights reserved. MPLS v2.1—6-15

Introducing Managed CE Routers Service

Overview

A service provider can use a separate network management VPN to manage the CE routers of all the VPNs. A pair of RT extended communities is used to accomplish this goal. One RT exports CE router loopback addresses and is imported into the VRF of the network management VPN. The other RT exports the networks from the VRF associated with the network management VPN and imports them into all other VRFs. This lesson explains some of the requirements and the implementation solution for the managed CE routers service.

It is important to be able to recognize the requirements of the network and to match them with customer requests. This lesson looks at one such requirement and explains how to handle it.


Objectives

Upon completing this lesson, you will be able to identify the characteristics of the managed CE routers service. This ability includes being able to meet these objectives:

- Identify the overall requirements of a managed CE routers VPN
- Identify the VRF and RD requirements of a managed CE routers VPN
- Describe how to configure a managed CE routers VPN

What Are the Requirements of Managed CE Routers?

This topic identifies the overall requirements of a managed CE routers VPN.

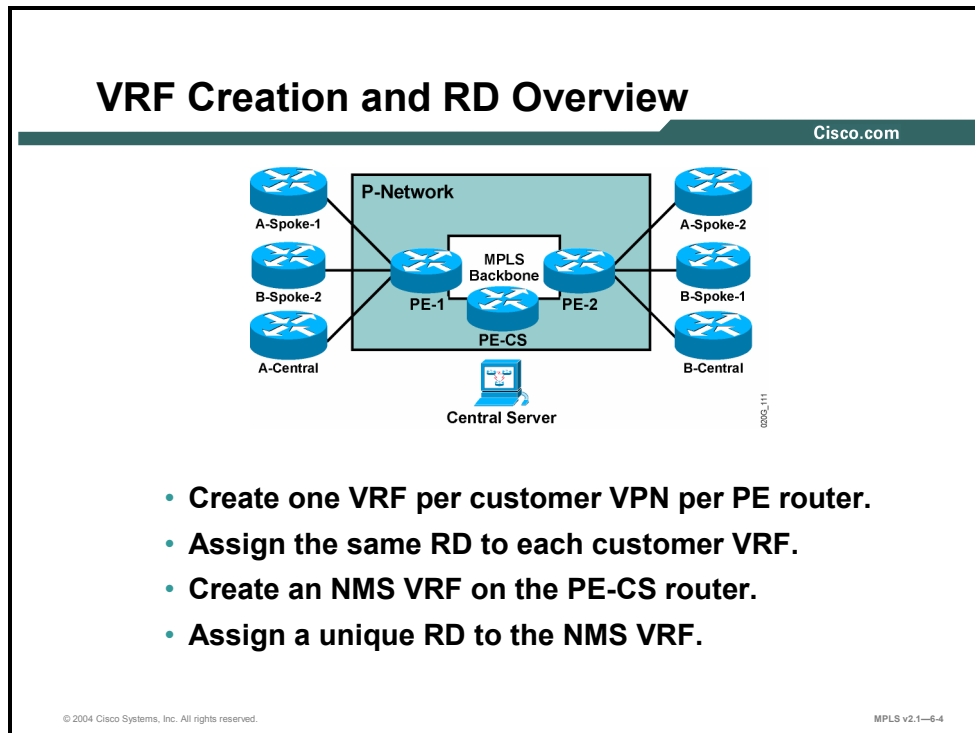


If the service provider is managing the customer routers, it is convenient to have a central point that has access to all CE routers but does not have access to the other destinations at the customer sites. This requirement is usually implemented by deploying a separate VPN for management purposes. This VPN needs to see all the loopback interfaces of all the CE routers. All CE routers have 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.

Note The topology described in this lesson is sometimes referred to as a “gray” network management VPN implementation, because all CE routers are accessed through a single link between the network management system (NMS) CE router and the network core. An alternative solution (a “rainbow” network management VPN), in which the NMS CE router has separate connections to each managed CE router, is usually used in combination with overlay VPNs (for example, Frame Relay networks).

What Are the VRF and RD Requirements?

This topic identifies the VRF and RD requirements of a managed CE routers VPN.



The VRF and RD design is the same as with central services VPNs. The only difference between this topology and the central services VPN topology combined with a simple VPN topology is the RT marking process during route export.

Configuring Managed CE Routers

This topic describes how to configure a managed CE routers VPN.

Configuring Route Targets

Cisco.com

- Configure the per-customer import-export route target in all customer VRFs.
- Configure the NMS import-export route target in NMS VRF.
- Import routes with the NMS RT into the customer VRF.
- Export loopback addresses from customer VRF with RT NMS_Client.
- Import routes with RT NMS_Client into NMS VRF.

© 2004 Cisco Systems, Inc. All rights reserved. MPLS v2.1—6-5

This table shows an RD and RT numbering scheme for PE-1.

PE-1 RD and RT Numbering Scheme


VRF	RD	Import RT	Export RT
VPN-A	123:750	123:750 123:101	123:750 123:100 (NMS_Client)
VPN-B	123:760	123:760 123:101	123:760 123:100 (NMS_Client)

This table shows an RD and RT numbering scheme for PE-CS.

PE-CS RD and RT Numbering Scheme

VRF	RD	Import RT	Export RT
NMS	123:101	123:101 123:100 (NMS_Client)	123:101

Configuring VRFs



You can have a configuration for a customer VRF with differentiated RT export for loopback addresses.

Example: Configuring VRFs

The figure illustrates this example. An export route map is used to match one part of the IP address space and attach an additional RT to the routes within this address space (CE router loopback addresses).

Note The routing protocol between PE and CE routers has to be secured (with distribute lists or prefix lists) to prevent customers from announcing routes in the address space dedicated to network management; otherwise, customers can gain two-way connectivity to the network management station.

The CE router loopback addresses are then imported into the server VPN based on the additional RT attached to them during the export process.

Note This design allows client sites to send packets to the network management VPN regardless of the source address. Special precautions should be taken to protect the network management VPN from potential threats and denial-of-service attacks coming from customer sites.

Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- **The managed CE routers service allows the service provider to access the loopback addresses of the CE router for management purposes.**
- **Managed VRF and RD design is the same as with the hybrid of a central and a simple VPN.**
- **Managed RT design is the same as with the hybrid of central and simple VPN, except for the RT marking process during route export.**

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—6-7

Introducing MPLS Managed Services

Overview

Market forces are driving service providers to provide additional centralized services to their customers. In addition, these services need to be integrated with existing VPN service. To meet this need, Cisco has provided a set of VPN-aware managed services. This lesson discusses Cisco MPLS managed services, focusing on which service provider needs can be met by Cisco MPLS managed services, and how those services are implemented in an MPLS network.

To successfully implement managed services, you need to understand the needs of the service provider, which kind of service can meet those needs, and how that service is implemented.

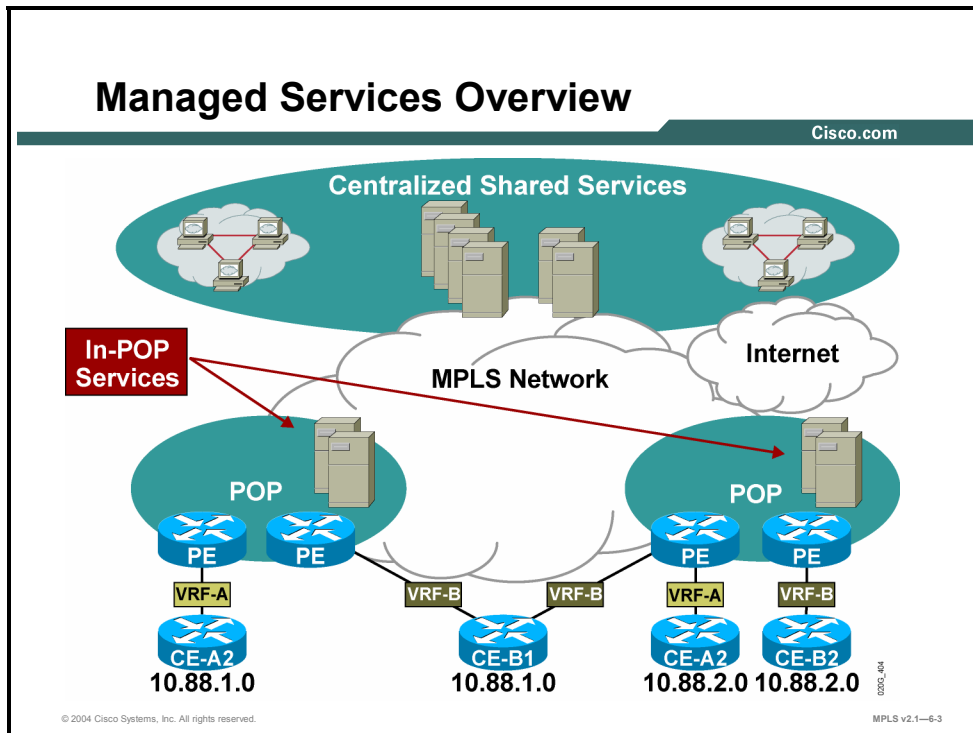
Objectives

Upon completing this lesson, you will be able to describe how Cisco MPLS VPN managed services can be implemented. This ability includes being able to meet these objectives:

- Describe the features of Cisco MPLS VPN managed services
- Describe the features of managed NAT
- Describe the features of managed DHCP relay
- Describe the features of managed on-demand address pools
- Describe the features of managed HSRP and VRRP
- Describe the features of managed multicast VPNs

What Are MPLS VPN Managed Services?

This topic describes the features of Cisco MPLS VPN managed services.



In modern networks, many end users have a need to connect to common services, such as e-mail, DHCP servers, and so on. Typically, these services have been provided by individual enterprises as part of their network.

Cisco MPLS for Managed Shared Services is a set of features delivered in Cisco IOS software for enabling managed shared services for MPLS VPNs. Building on leading Cisco MPLS capabilities, service providers now can provide their enterprise clients all the connectivity benefits associated with Cisco MPLS VPNs while creating additional revenue streams by also providing economically attractive IP services.

Managed Services Overview (Cont.)

Cisco.com

- **NAT**
- **DHCP relay for MPLS VPNs**
- **ODAP for MPLS VPNs**
- **HSRP for MPLS VPNs**
- **VRRP for MPLS VPNs**
- **Multicast VPNs**

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—6-4

Cisco has expanded its widely deployed MPLS VPN solution to include the following technologies in Cisco IOS software:

- NAT for MPLS VPNs
- Dynamic Host Configuration Protocol (DHCP) relay for MPLS VPNs
- On-demand address pools (ODAP) for MPLS VPNs
- Hot Standby Router Protocol (HSRP) for MPLS VPNs
- Virtual Router Redundancy Protocol (VRRP) for MPLS VPNs
- Multicast VPNs

With these key new technologies, enterprise IP services can now be moved from the enterprise network into the MPLS VPN network of the service provider and shared across multiple VPNs for greater operational leverage and economies of scale.

Managed Services Overview (Cont.)

Cisco.com

- **Cisco MPLS for Managed Shared Services can eliminate the following problems commonly associated with delivering advanced services to MPLS VPN customers:**
 - Inefficiency in resource utilization
 - High traffic loads
 - Management complexity
- **Cisco MPLS technology incorporates features for:**
 - More effectively managing shared IP services
 - Delivering multicast-based services
 - Adding flexibility to client service selection

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—6-5

Cisco MPLS for Managed Shared Services eliminates many of the problems—such as inefficiency in resource utilization, high traffic loads, and management complexity—commonly associated with delivering advanced services to MPLS VPN customers. The Cisco MPLS technology incorporates features for more effectively managing shared IP services, delivering multicast-based services, and adding flexibility to client service selection.


The Cisco MPLS offering includes a number of VRF features that present opportunities for new IP services revenue streams and for cost savings. NAT for MPLS VPNs, for instance, allows service providers to more cost-effectively support services such as content hosting, enterprise resource planning (ERP) application hosting, and managed Internet access. Other features add support in the MPLS network for industry-standard protocols and improve or automate routing control. The comprehensive collection of functions can help service providers eliminate many customer-expressed barriers to entry by ensuring that MPLS VPN business clients have access to the robust functionality that they expect in the enterprise environment.

Cisco MPLS for Managed Shared Services also incorporates multicast VPN functionality to help service providers meet enterprise market demands for IP services that are essential in applications such as telecommuting. By reducing packet replication in the MPLS network, multicast VPN technology allows for massively scalable distribution of data, voice, and video streams. Utilizing multicast VPN features, service providers can leverage existing infrastructure resources to offer competitive services in videoconferencing, e-learning, and other Internet-based streaming applications.

Taken together, the Cisco MPLS for Managed Shared Services features give service providers powerful new MPLS VPN functionality and versatility—without deployment or management complexity.

What Is Network Address Translation?

This topic describes the features of NAT services in an MPLS VPN environment.




In modern MPLS networks, enterprises have to pay for leased links and router ports for Internet connectivity in addition to VPN connectivity, and also the operational expenses associated with internally managing NAT. While service providers can currently provide NAT services to their enterprise clients with additional router and NAT devices, it is a highly complex design. NAT for MPLS VPNs is a simpler and more flexible way to integrate NAT services within MPLS VPNs with a single network connection that provides both MPLS VPN connectivity and access to shared services.

Because NAT for MPLS VPNs provides more economical NAT services, these services can be made more appealing to enterprise clients with a resulting revenue opportunity for service providers.

Network Address Translation: NAT for Shared Services

Cisco.com



- **Reduced capital and operating expenditures for the enterprise:**
 - Leased lines and router ports for Internet connectivity go away
 - Reduces network complexity; reduces operations and manageability cost
- **Increased revenues for service providers:**
 - Can provide outsourced NAT services to enterprises

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—6-7

The integration of NAT with MPLS VPNs allows multiple MPLS VPNs to be configured on a single device to work together. NAT can differentiate from which MPLS VPN that it receives IP traffic, even if the MPLS VPNs are all using the same IP addressing scheme. This enhancement enables multiple MPLS VPN customers to share services while ensuring that each MPLS VPN is completely separate from the other.

MPLS service providers would like to provide value-added services such as Internet connectivity, Domain Name System (DNS) servers, and Voice over IP (VoIP) to their customers. These additions require that customer IP addresses be different than the service provider services. Because MPLS VPNs allow customers to use overlapped IP addresses in their networks, NAT must be implemented to make the services possible.

Cisco NAT for MPLS VPNs provides the following:

- A simple and more flexible way of integrating NAT with MPLS VPNs
- Automatic management of the overlapping of VPN address spaces (allowable in MPLS VPNs) to ensure that addresses are mapped correctly in shared-services applications
- Centralized delivery of full-VPN NAT services
- NAT redundancy (NAT can be configured on one or more PE routers.)

Cisco NAT for MPLS VPNs eliminates the requirement for physical connectivity between a shared service and the provider network that is performing NAT.

Network Address Translation: Implementation

Cisco.com

- The inside interface can be any type of interface (both MPLS and non-MPLS).
- Outside interface:
 - The outside interface can be part of a VRF or a regular “generic” physical or logical interface.
 - MPLS label switching cannot be enabled on these interfaces.
- NAT can be configured on one or more PE routers for redundancy:
 - The “shared service” does not need to be physically connected to the PE device performing NAT.
 - NAT pools must be unique.
- NAT will inspect all traffic routed VRF to VRF or VRF to global routing table.
- VPN-aware NAT maintains support for all existing applications and protocols in an MPLS VPN environment.

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—6-8

NAT can be implemented on the PE router in the following scenarios:

- **Service point:** Shared access can be from a generic interface or from a VPN interface.
- **NAT point:** NAT can be configured on the PE router that is directly connected to the shared access gateway or on the PE router that is not directly connected to the shared access gateway.
- **NAT interface:** The shared access gateway interface is most often configured as the outside interface of NAT. The inside interface of NAT can be either the provider edge-customer edge (PE-CE) interface of a VPN, the interface to the MPLS backbone, or both. The shared access gateway interface can also be configured as the inside interface.
- **Routing type:** The common service can be Internet connectivity or a common server. For Internet connectivity, a default route should be propagated to all the VPN customers that use the service. For common server access, a static or dynamically learned route should be propagated to the VPN customers.
- **NAT configuration:** NAT can have the following different configurations: static, dynamic, pool or interface overloading, and route map.

The concept of VPN-aware NAT is very similar to classic NAT. Inside and outside interfaces serve the same function as in classic NAT; only the location of the NAT service is changed. An inside interface can be any type of interface. An outside interface can be part of a VRF or a regular “generic” physical or logical interface, but MPLS cannot be enabled on these interfaces.


The “shared service” does not need to be physically connected to the PE device performing NAT. In addition, NAT can be configured on one or more PE routers for redundancy.

NAT will inspect all traffic routed VRF to VRF or VRF to global routing table to determine when and where NAT should be applied.

VPN-aware NAT also maintains support for all existing applications and protocols in an MPLS VPN environment.

Network Address Translation: Implementation (Cont.)

Cisco.com



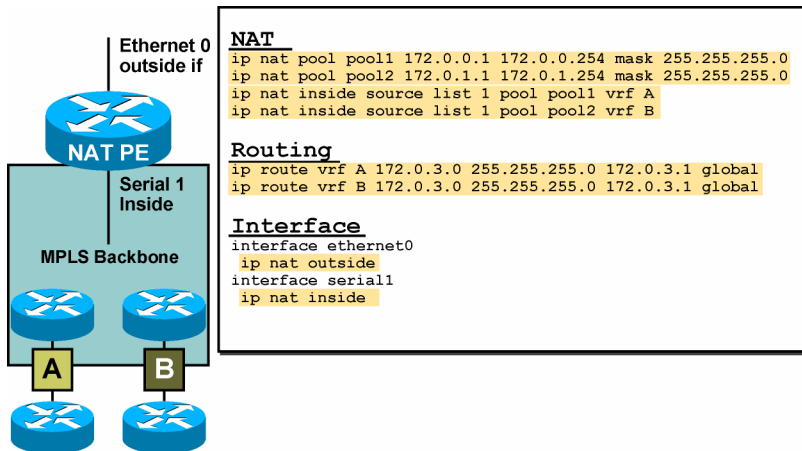
There are two approaches to implementing NAT in the MPLS VPN network. NAT can be implemented on the CE router, which is already supported by NAT, or it can be implemented on a PE router. Integration of NAT with MPLS VPNs enables the implementation of NAT on a PE router in an MPLS cloud.

Example: Network Address Translation

The figure presents an example of VPN-aware NAT. CE-A1, CE-A2, CE-B1, and CE-B2 are clients in VRF-A and VRF-B. Packets from these clients destined for the shared service are routed to the inside interface of the NAT PE router over their respective VPNs. At the NAT PE router, the address translation process replaces the inside source address with the outside source address from the NAT table and forwards the packet to the shared service.

Network Address Translation: Implementation with Multiple NAT pools

Cisco.com



NAT services can be configured on any router that is part of the VPN.

Example: NAT Implementation with Multiple NAT Pools

The figure presents an example of VPN-aware NAT configuration for two VPNs, A and B. NAT services are being configured on the PE router connected to the shared services.

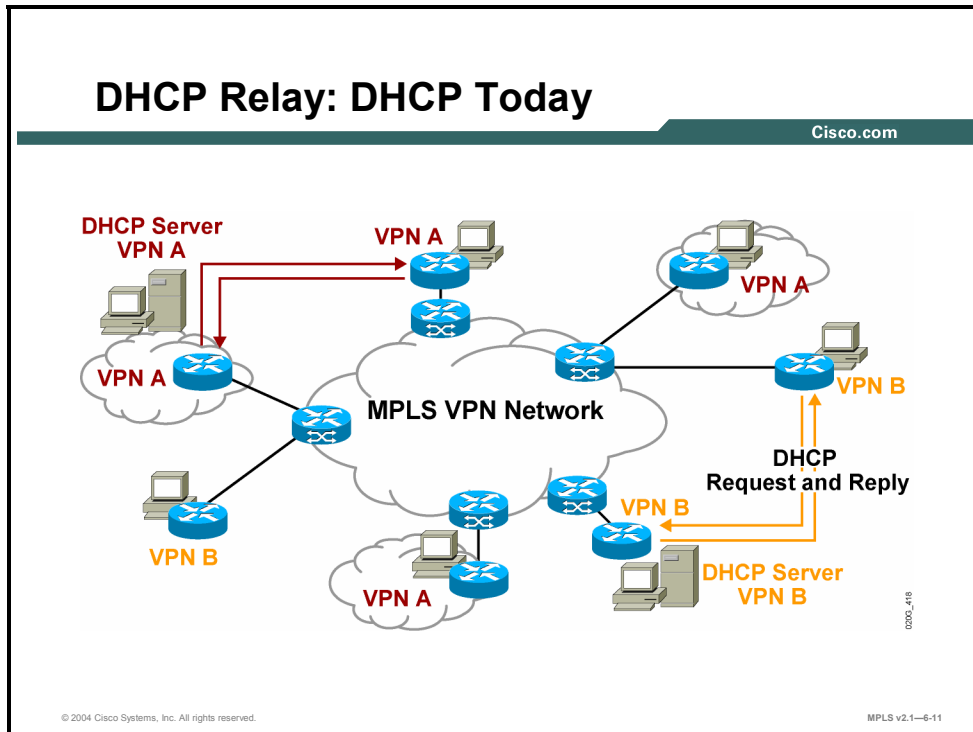
NAT pools are configured with a standard NAT configuration command: **ip nat pool**. Only one NAT pool is required; however, in this example, there are two pools, one for each VPN to allow for easy address administration. The NAT pools are assigned to their respective VPNs using the **ip nat inside pool** command.

NAT services are applied to the interfaces using the **ip nat** command.

Because the outside interface is not participating in the VPN, the VPN VRF would not normally know of its existence. A static default route is created, pointing to the next-hop address of the shared services for each VPN, by using the **ip route vrf** command with the **global** keyword.

What Is DHCP Relay?

This topic describes the features of managed DHCP relay services in an MPLS VPN environment.

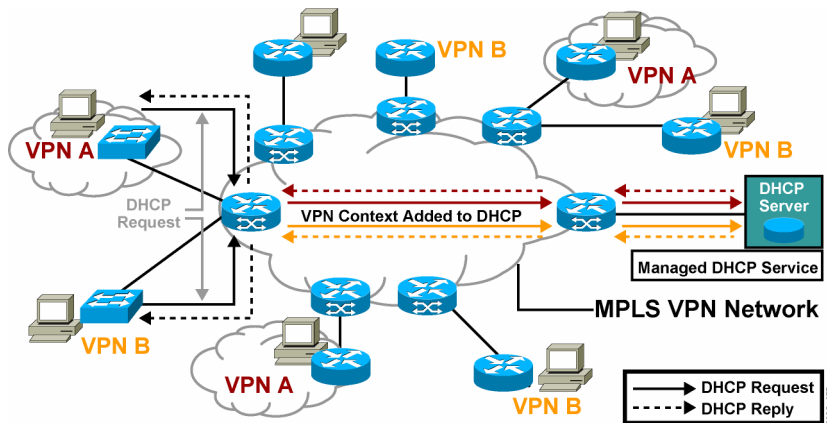


Current implementations of DHCP suffer from the following issues:

- Even if the DHCP servers are collocated, there is a replication of DHCP servers per VPN.
- There is no added value from the service provider.

DHCP Relay: DHCP Support for Shared Services

Cisco.com



Service providers can take advantage of another centralized service to support DHCP clients. DHCP Relay for MPLS VPNs enables a DHCP relay agent to forward information about the DHCP request and the VPN association when forwarding client-originated DHCP packets to a DHCP server. The DHCP server can then use that information to interpret IP addresses or other parameter-assignment policies. The DHCP relay agent information option is organized as a single DHCP option that contains one or more suboptions (VPN identifier, subnetwork selection, and server identifier override) to convey information known by the relay agent.

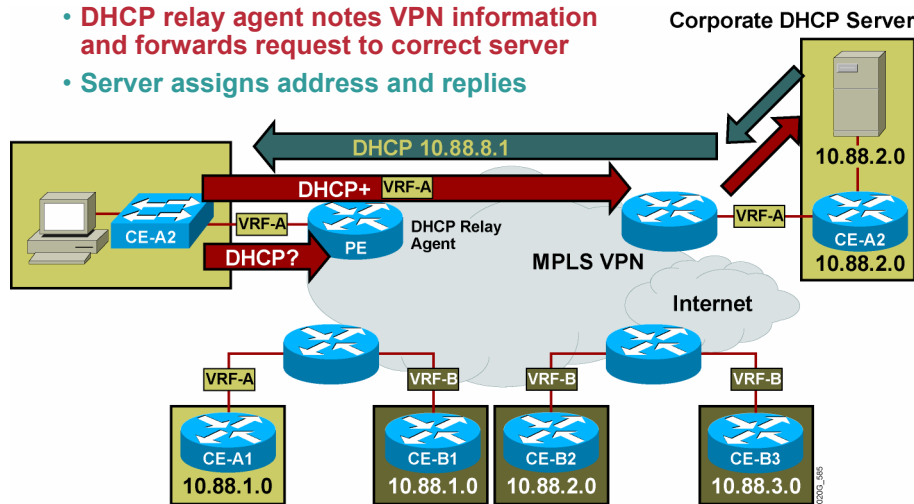
The DHCP relay agent information option (option 82) enables a DHCP relay agent to include information about itself when forwarding client-originated DHCP packets to a DHCP server. The DHCP server can use this information to implement IP address or other parameter-assignment policies. The DHCP relay agent information option is organized as a single DHCP option that contains one or more suboptions that convey information known by the relay agent.

In some environments, a relay agent resides in a network element that also has access to one or more MPLS VPNs. A DHCP server that wants to offer service to DHCP clients on those different VPNs needs to know the VPN in which each client resides. The network element that contains the relay agent typically knows about the VPN association of the DHCP client and includes this information in the relay agent information option.

DHCP Relay: Corporate DHCP Server

Cisco.com

- End station makes DHCP request
- DHCP relay agent notes VPN information and forwards request to correct server
- Server assigns address and replies



© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—6-13

Example: DHCP Relay—Corporate DHCP Server

In this two-VPN example, a corporate DHCP server and a DHCP client have been added to VPN-A. The client broadcasts a DHCP request to the local relay. The local relay converts the broadcast to a unicast request for the DHCP server and adds the VPN ID. This request is forwarded to the egress PE router based upon the DHCP server address. From the egress PE router, the request is forwarded to the DHCP server.

The DHCP server assigns the client an address and replies to the DHCP relay, which in turn forwards the reply to the client.

The relay agent uses the VPN identifier suboption to tell the DHCP server the VPN for every DHCP request that it passes on to the DHCP server. This suboption is also used to properly forward any DHCP reply that the DHCP server sends back to the relay agent. The VPN identifier suboption contains the VPN ID configured on the incoming interface to which the client is connected. If you configure the VRF name but not the VPN ID, the VRF name is used as the VPN identifier suboption. If the interface is in global routing space, the VPN suboptions are not added.

The subnetwork selection suboption allows the separation of the subnetwork where the client resides from the IP address used to communicate with the relay agent. In typical DHCP processing, the gateway address specifies both the subnetwork on which a DHCP client resides and the IP address that the server can use to communicate with the relay agent. Situations exist where the relay agent needs to specify a subnetwork on which a DHCP client resides that is different from the IP address that the server can use to communicate with the relay agent. The subnetwork selection suboption is included in the relay agent information option and passed on to the DHCP server. The gateway address is changed to the outgoing interface of the relay agent toward the DHCP server. The DHCP server uses this gateway address to send reply packets back to the relay agent.


The server identifier override suboption value is copied in the reply packet from the DHCP server instead of the normal server ID address. The server identifier override suboption contains the incoming interface IP address, which is the IP address on the relay agent that is accessible from the client. Using this information, the DHCP client sends all renew and release packets to the relay agent. The relay agent adds all of the VPN suboptions and then forwards the renew and release packets to the original DHCP server.

After the relay agent has added these suboptions to the DHCP relay agent information option, the gateway address is changed to the outgoing interface of the relay agent toward the DHCP server. When the packets are returned from the DHCP server, the relay agent removes the relay agent information options and forwards the packets to the DHCP client on the correct VPN.

DHCP Relay: Shared DHCP Server

Cisco.com

- End station makes DHCP request
- DHCP relay agent adds VPN information at PE router (cable or DSL headend Layer 2 access)
- Server assigns address based on option 82 data and replies



Typical network topology for a shared DHCP server involves a bridged access to the remote location via DSL or cable modem technology. Aggregation into the PE router involves the use of headend equipment, such as digital subscriber line access multiplexer (DSLAM) and cable modem termination system (CMTS).

Example: DHCP Relay—Shared DHCP Server


In this figure, the corporate DHCP server has been replaced with a shared DHCP server provided by the service provider. Because the server is shared between VPNs, a NAT PE router could also be included to provide address translation.

The client broadcasts a DHCP request to the local relay. The local relay converts the broadcast to a unicast request for the shared DHCP server and adds the VPN ID. This request is forwarded to the egress PE router via the NAT PE router based upon the DHCP server address. At the NAT PE router, an address translation is performed and the request is forwarded to the DHCP server.

The DHCP server assigns the client an address from the VPN pool and replies to the DHCP relay, which in turn forwards the reply to the client.

What Are On-Demand Address Pools?

This topic describes the features of managed ODAP services in an MPLS VPN environment.




Today, service providers face the following challenges concerning efficient management of IP address space for customers:

- Address management is independent but inefficient.
- Providers need to manage addresses manually and allocate them to RADIUS or DHCP servers.
- Once site thresholds are reached, new addresses have to be manually allocated.

On-Demand Address Pools: Shared Service

Cisco.com



- Use IOS DHCP server
- Request address pools on demand when a threshold is reached
- Much more efficient IP address management for MPLS VPNs
- Less network load and new revenue opportunity for service provider

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—6-17


With MPLS VPNs, service providers have to allocate their IP address pools to independent RADIUS or DHCP servers for each VPN. Once the site threshold has been reached, new addresses have to be provided manually. With ODAP for MPLS VPNs, this process can now be fully automated and provided as a shared service on one or more servers. When the site threshold is exceeded, ODAP automates the process of expanding the overall address pool, reducing network loading and performing configuration.

The Cisco ODAP for MPLS VPNs feature provides the following:

- Capabilities for automated control
- Support for MPLS VPNs, with addresses assigned per subnetwork, per interface
- Easy monitoring capabilities (Pool manager can assess address utilization and expand the pool as needed.)
- Simplified VPN setup (Upon configuration, pool manager can request an initial subnet from the address pool server.)

On-Demand Address Pools: Provisioning and Startup

Cisco.com



- The PE router is configured for ODAP.
- **ODAP requests initial pool for VRF-A from the server.**
- The server replies with the initial address pool.


© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—6-18

As soon as the DHCP server is enabled on the PE router that has ODAP enabled, the PE router will request an initial address pool from its designated server.

On-Demand Address Pools: Address Pool Management

Cisco.com



- The end station makes the DHCP request.
- **The DHCP server fulfills request from pool—reaches 90 percent.**
- The ODAP pool manager requests expansion.
- The server allocates another subnetworks and replies.
- The PE router adds subnetworks routing information to the VRF.

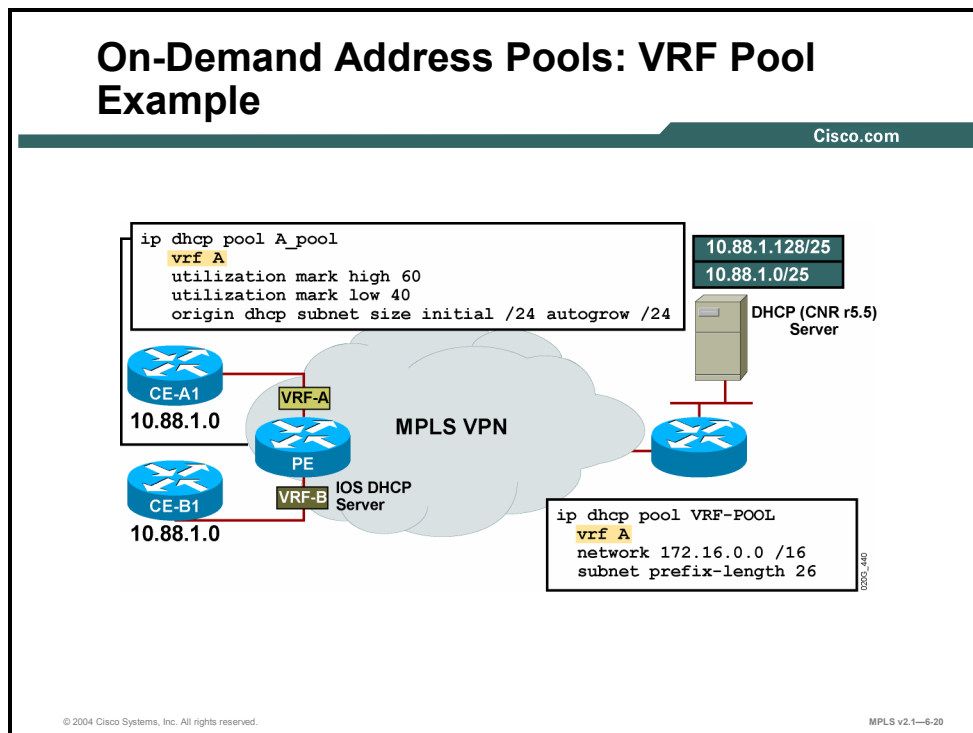
© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—6-19

The PE router will honor DHCP requests and assign addresses until its address pool is 90 percent depleted. At this point, the PE router will request an extension of the address pool from its designated server.

Example: On-Demand Address Pools

In this example, the PE router is configured as a DHCP server with ODAP, and a second router is configured to be a subnetwork allocation server for VNP A.



On the PE router, a DHCP pool named “A_pool” has been created. This pool is associated with VPN A. Three new commands have been introduced to the DHCP command presented in the previous topic. To configure an address pool as an ODAP, use the **origin** command. The **subnet size initial** size option is used to set the initial size of the first requested subnetwork. You can enter *size* as either the subnet mask (nnnn.nnnn.nnnn.nnnn) or prefix size (/nn). The **autogrow** size option is used to specify that the pool can grow incrementally. The *size* argument is the size of the requested subnetworks when the pool requests additional subnetworks (upon detection of high utilization). You can enter *size* as either the subnet mask (nnnn.nnnn.nnnn.nnnn) or prefix size (/nn).


To configure the high utilization mark of the current address pool size, use the **utilization mark high** command in DHCP pool configuration mode. To configure the low utilization mark of the current address pool size, use the **utilization mark low** command in DHCP pool configuration mode.

On the ODAP server, a VRF subnetwork allocation pool named “VRF-POOL,” which allocates subnetworks from the 172.16.0.0/16 network, has been configured to match the VRF named “A.” The configuration of the **subnet prefix-length** command in this example configures the size of each subnetwork that is allocated from the subnetwork pool to support 62-host IP addresses.

Note Additional DHCP and ODAP configuration options are available. Refer to the Cisco IOS documentation for further information.

What Are HSRP and VRRP?

This topic describes the features of managed HSRP and VRRP services in an MPLS environment.



Today, service providers face challenges when it comes to implementing an efficient redundancy scheme. To address many of these issues, Cisco has offered HSRP and VRRP, which have provided the following benefits:

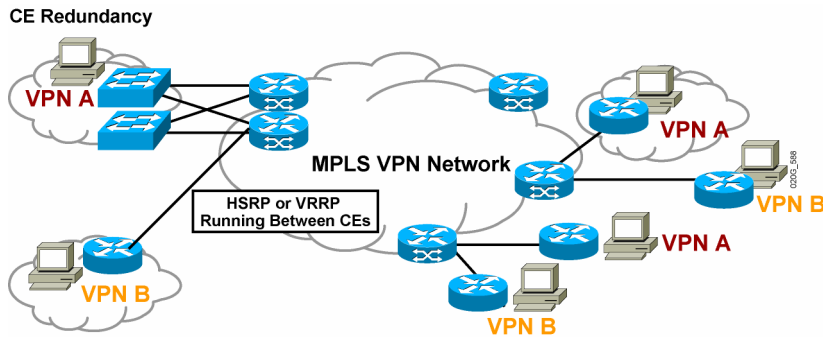
- Both protocols protect against first-hop router failure.
- Hosts access a virtual IP address that represents a cluster of gateway routers.
- If a gateway router fails, others in the cluster automatically take over.

However, when implemented in a VPN environment, HSRP and VRRP have these limitations:

- They provide CE redundancy only.
- Dual homing to PEs can be done but does not provide fast failover.

HSRP and VRRP Today: Support for MPLS VPNs

Cisco.com



- Improved network availability
- HSRP:
 - Transparent network topology modifications
 - Simple, centralized control of hot standby parameters
- VRRP:
 - Standards-based protocol
 - The flexibility to choose the protocol that best suits each environment

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—6-22

Cisco MPLS for Managed Shared Services also provides HSRP support on MPLS VPN interfaces. This feature provides transparent “first-hop IP routing” redundancy for workstations or routers connected to interfaces within the MPLS VPN. HSRP creates a hot standby router group with a lead router that services all packets sent to the hot standby address. Other routers within the group monitor the lead router. If the lead fails, a standby router inherits the lead position and also the hot standby address. The HSRP protocol allows specification of active routers, preemption delays, hold times, and interface status tracking.

The benefits of HSRP for MPLS VPNs include the following:

- Improved network availability
- Transparent network topology modifications
- Simple, centralized control of hot standby parameters

Similar to HSRP, VRRP allows a group of routers to function as one virtual router. Cisco MPLS for Managed Shared Services includes VRRP for MPLS VPNs by enabling the group of routers to share one virtual IP address and one virtual MAC address. One master router performs packet forwarding for the local hosts, and the rest of the routers within the group can act as backup routers to protect from failures of the master. With VRRP, the backup routers stay idle as far as packet forwarding is concerned.

The benefits of VRRP for MPLS VPNs include the following:

- Improved network availability
- Standards-based protocol
- The flexibility to choose the protocol that best suits each environment

HSRP and VRRP Today: VPN A HSRP Example


Cisco.com

Router PE1 Configuration

```
!
ip vrf A
rd 100:1
route-target export 100:1
route-target import 100:1
!
interface ethernet0
ip vrf forwarding A
ip address 10.2.0.1 255.255.0.0
standby 1 ip 10.2.0.20
standby 1 priority 105 preempt delay 10
standby 1 timers 1 3
standby 1 track ethernet1 10
standby 1 track ethernet2 10
```

Router PE2 Configuration

```
!
ip vrf A
rd 100:1
route-target export 100:1
route-target import 100:1
!
interface ethernet0
ip vrf forwarding A
ip address 10.2.0.2 255.255.0.0
standby 1 ip 10.2.0.20
standby 1 priority 100 preempt delay 10
standby 1 timers 1 3
standby 1 track ethernet1 10
standby 1 track ethernet2 10
```



© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—6-23

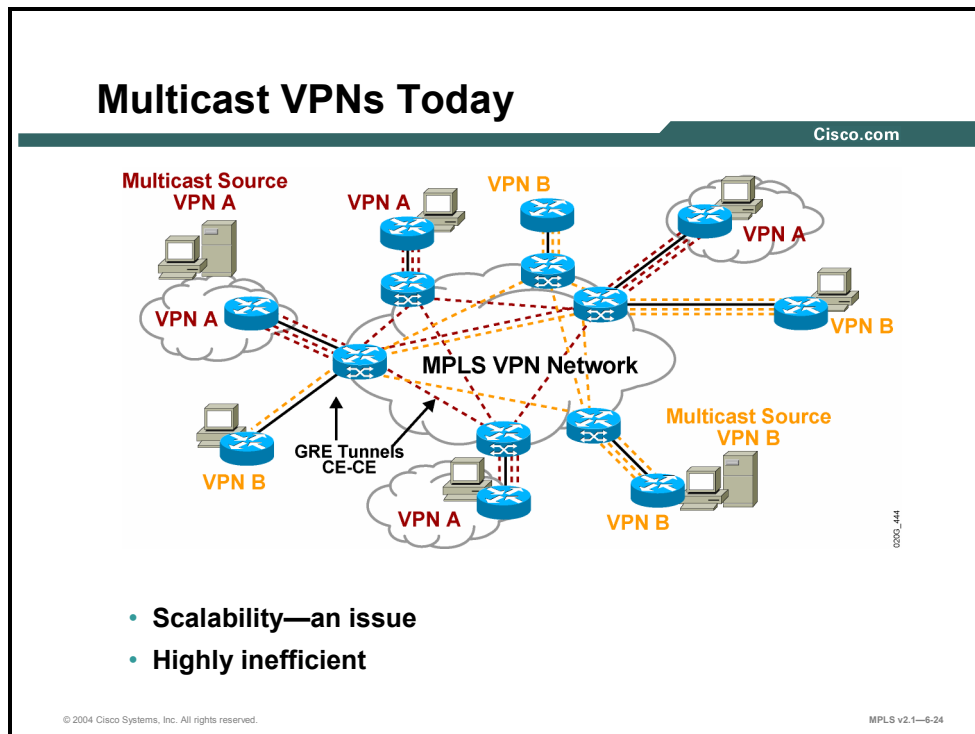
The creation of a VPN-aware HSRP is a combination of the standard MPLS VPN and HSRP commands.

Example: HSRP and VRRP Today

In this figure, a VPN-aware HSRP is created for VPN A. MPLS forwarding has been enabled on Ethernet0, and the virtual IP address has been configured as 10.2.0.20.

What Are Multicast VPNs?

This topic describes the features of managed multicast VPN services in an MPLS VPN environment.



Historically, IPinIP generic routing encapsulation (GRE) tunnels were the only way to connect through a service provider network. Although such tunneled networks tend to have scalability issues, they represented the only means of passing IP multicast traffic through a VPN.

MPLS was derived from tag switching, and various other vendor methods of IP switching support enhancements in the scalability and performance of IP-routed networks by combining the intelligence of routing with the high performance of switching. MPLS is now used for VPNs, which is an appropriate combination because MPLS decouples information used for forwarding of the IP packet (the label) from the information carried in the IP header.


A multicast VPN allows an enterprise to transparently interconnect its private network across the network backbone of a service provider. The use of a multicast VPN to interconnect an enterprise network in this way does not change the way that the enterprise network is administered and it does not change general enterprise connectivity.

The multicast VPN feature in Cisco IOS software provides the ability to support the multicast feature over a Layer 3 VPN. As enterprises extend the reach of their multicast applications, service providers can accommodate these enterprises over their MPLS core network. IP multicast is used to stream video, voice, and data to an MPLS VPN network core.

A VPN represents network connectivity across a shared infrastructure, such as an Internet service provider (ISP). Its function is to provide the same policies and performance as a private network, at a reduced cost of ownership, thus creating many opportunities for cost savings through operations and infrastructure.

Multicast VPNs Today (Cont.)

Cisco.com



- Enabling service providers with MPLS VPN networks to offer multicast services to their enterprise clients
- Minimizing configuration time and complexity (Configuration is required only at edge routers.)
- Ensuring transparency of the service provider network
- Providing the ability to easily build advanced enterprise-friendly services, such as virtual multicast networks
- Increasing network scalability

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—6-25

By implementing native multicast functionality inside their MPLS VPN networks, service providers can now monetize multicast services. Service providers can utilize current resources to support bandwidth-hungry streaming services, such as telecommuting, videoconferencing, e-learning, and a host of other business applications. Cisco multicast VPN technology helps improve the efficiency of the bandwidth-hungry applications of enterprise networks by eliminating the packet replication and performance issues associated with the distribution of multicast traffic.

Multicast VPNs benefit service providers by accomplishing the following:

- Enabling service providers with MPLS VPN networks to offer multicast services to their enterprise clients
- Minimizing configuration time and complexity (Configuration is required only at edge routers.)
- Ensuring transparency of the service provider network
- Providing the ability to easily build advanced enterprise-friendly services, such as virtual multicast networks
- Increasing network scalability

Multicast VPNs Today: Terminology

Cisco.com

- **mVPN**: VPN that supports multicast natively
- **MVRF**: VRF that supports unicast and multicast tables
- **MDT**: A multicast tree, built in the core network between PE and P routers, that distributes multicast traffic between sites
- **Default MDT**: Default MDT group used for control traffic and flooding channel for dense-mode and low-bandwidth groups.
- **Data MDT**: MDT group created on demand for mVPN (S,G) pairs—usually high-bandwidth traffic

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—6-26

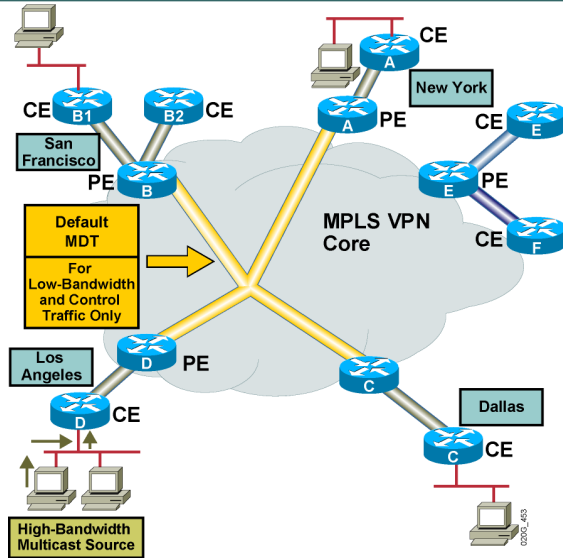
VPN-aware multicast technology has introduced a new set of terminology.

Multicast VPNs introduce multicast routing information to the VRF table. When a PE router receives multicast data or control packets from a CE router, forwarding is performed according to information in the multicast virtual routing and forwarding instance (MVRF).

A set of MVRFs that can send multicast traffic to each other constitutes a multicast domain. For example, the multicast domain for a customer that wanted to send certain types of multicast traffic to all global employees would consist of all CE routers associated with that enterprise.

Multicast VPNs Today: Default MDT

Cisco.com



- Customer CE devices join the MPLS core through provider PE devices.
- The MPLS core forms a default MDT for a given customer.

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—6-27

Multicast VPNs establish a static default multicast distribution tree (MDT) for each multicast domain. The default MDT defines the path used by PE routers to send multicast data and control messages to every other PE router in the multicast domain.

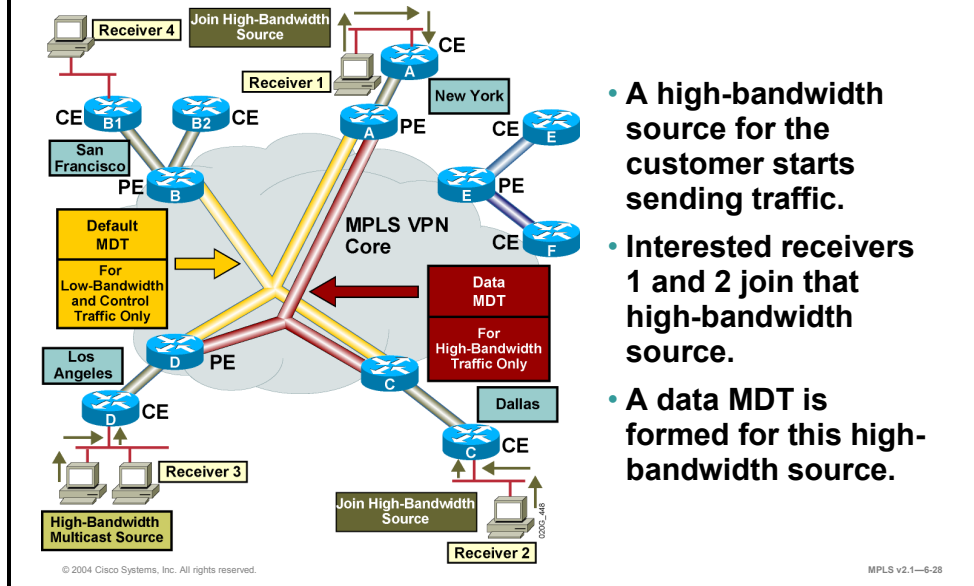
Example: Multicast VPNs—Default MDT

In the example, a service provider has a multicast customer with offices in San Francisco, Los Angeles, New York, and Dallas. A one-way multicast presentation is occurring in Los Angeles. The service provider network supports all three sites associated with this customer, in addition to the site of a different enterprise customer.

The default MDT for the enterprise customer consists of provider routers and their associated PE routers. The other PE router is not part of the default MDT, because it is associated with a different customer.

Multicast VPNs Today: Data MDT

Cisco.com



- A high-bandwidth source for the customer starts sending traffic.
- Interested receivers 1 and 2 join that high-bandwidth source.
- A data MDT is formed for this high-bandwidth source.

Multicast VPNs also support the dynamic creation of MDTs for high-bandwidth transmission. Data MDTs are a feature unique to Cisco IOS software. Data MDTs are intended for high-bandwidth sources, such as full-motion video inside the VPN, to ensure optimal traffic forwarding in the MPLS VPN core. The threshold at which the data MDT is created can be configured on a per-router or a per-VRF basis.

When the multicast transmission exceeds the defined threshold, the sending PE router creates the data MDT and sends a User Datagram Protocol (UDP) message that contains information about the data MDT to all routers in the default MDT. The statistics to determine whether a multicast stream has exceeded the data MDT threshold are examined once every 10 seconds. If multicast distributed switching is configured, the time period can be up to twice as long.

Data MDTs are created only for (S, G) multicast route entries within the VRF multicast routing table. Data MDTs are not created for (*, G) entries regardless of the value of the individual source data rate.

Example: Multicast VPNs—Data MDT

In the example here, an employee joins the multicast session. The PE router associated with the employee site sends a join request that flows across the default MDT for the multicast domain of the customer. The PE router associated with the multicast session source receives the request.


The source CE router begins to send the multicast data to the associated PE router, which sends the multicast data along the default MDT. Immediately after sending the multicast data, the source PE router recognizes that the multicast data exceeds the bandwidth threshold at which a data MDT should be created. Therefore, the PE router creates a data MDT, sends a message to all routers using the default MDT that contains information about the data MDT, and, 3 seconds later, begins sending the multicast data for that particular stream using the data MDT. Only the PE routers that have interested receivers for this source will join the data MDT and receive traffic on it.

PE routers maintain a Protocol Independent Multicast (PIM) relationship with other PE routers over the default MDT, and a PIM relationship with their directly attached PE routers.

The figure depicts the final flow of multicast data sourced from the multicast sender in Los Angeles to the multicast clients in New York and Dallas. Multicast data sent from the multicast sender is delivered in its original format to its associated PE router using sparse mode, Bi-directional (Bi-Dir) or Source Specific Multicast (SSM). This PE router then encapsulates the multicast data and sends it across the data MDT using the configured MDT data groups. The mode used to deliver the multicast data across the data MDT is determined by the service provider and has no direct correlation with the mode used by the customer. The PE router in New York receives the data along the data MDT. That PE router de-encapsulates the packet and forwards it in its original format toward the multicast client using the mode configured by the customer.

Multicast VPNs Today: Solution Concept

Cisco.com



- P and PE must be routers multicast-enabled.
- Global multicast routing tables are created in the provider network.
- Globally, PE routers configured to run PIM (global instance) with adjacent P routers.
- Multicast-enabled VPNs have a VPN multicast routing table (MVRF).
- There is no requirement to run same multicast protocols in the customer and provider network.
- If PIM is configured, PE routers maintain PIM adjacencies with CE devices:
 - No PIM adjacency between CE devices not directly connected
- Normal PIM configuration in customer network:
 - RPs, RRs, and so on

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—6-29

For every multicast domain of which an MVRF is a part, the PE router creates a multicast tunnel interface. A multicast tunnel interface is an interface that the MVRF uses to access the multicast domain. It can be thought of as a conduit that connects an MVRF and the global MVRF. One tunnel interface is created per MVRF.

To support VPN-aware multicast systems, PIM multicast (PIM, SSM, and so on) capability must be enabled on all affected P and PE routers. This addition results in a global multicast routing table being created in the provider network routers. The PE routers that have been configured to run PIM (global instance) will establish a PIM adjacency with neighboring provider routers (P routers).


Multicast-enabled VPNs will create a VPN multicast routing table (MVRF).

There is no requirement to run the same multicast protocols in the customer and provider network. If PIM is configured as the CE-to-PE multicast protocol, the PE devices maintain PIM adjacencies with CE devices. No PIM adjacency will be established between CE devices that are not directly connected.

Normal PIM configuration—for example, rendezvous point (RP), Router Reflector (RR), confederations—is accomplished in the customer network.

Multicast VPNs Today: Configuration Example

Cisco.com



```
!
ip vrf A
rd 1:1
route-target export 1:1
route-target import 1:1
mdt default 232.0.0.1
mdt data 232.5.0.0 0.0.0.255 threshold 100

! ip multicast-routing
ip multicast-routing vrf A
!
interface FastEthernet1/0/0
ip vrf forwarding A
ip address 172.16.140.1 255.255.255.0
ip pim sparse-dense-mode
!
interface GigabitEthernet4/0/0
ip address 10.0.2.1 255.255.255.0
ip router isis
ip pim sparse-mode
ip route-cache distributed
tag-switching ip
!
```

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—6-30

Configuring VPN-aware multicast capability is a combination of standard VPN, standard multicast, and new VPN-aware multicast commands.

Enabling a VPN for Multicast Routing

This task enables a VPN for multicast routing.

PIM

PIM can operate in dense mode or sparse mode. It is possible for the router to handle both sparse groups and dense groups at the same time.

In dense mode, a router assumes that all other routers want to forward multicast packets for a group. If a router receives a multicast packet and has no directly connected members or PIM neighbors present, a prune message is sent back to the source. Subsequent multicast packets are not flooded to this router on the pruned branch. PIM builds source-based multicast distribution trees.

In sparse mode, a router assumes that other routers do not want to forward multicast packets for a group, unless there is an explicit request for the traffic. When hosts join a multicast group, the directly connected routers send PIM join messages toward the RP. The RP keeps track of multicast groups. Hosts that send multicast packets are registered with the RP by the first-hop router of that host. The RP then sends join messages toward the source. At this point, packets are forwarded on a shared distribution tree. If the multicast traffic from a specific source is sufficient, the first-hop router of the host may send join messages toward the source to build a source-based distribution tree.

Fast Switching and IP Multicast

Fast switching of IP multicast packets is enabled by default on all interfaces—including GRE and Distance Vector Multicast Routing Protocol (DVMRP) tunnels—with one exception: It is disabled and not supported over X.25 encapsulated interfaces. Note the following properties of fast switching:

- If fast switching is disabled on an incoming interface for a multicast routing table entry, the packet is sent at the process level for all interfaces in the outgoing interface list.
- If fast switching is disabled on an outgoing interface for a multicast routing table entry, the packet is process-switched for that interface, but may be fast-switched for other interfaces in the outgoing interface list.

Disable fast switching if you want to log debug messages, because when fast switching is enabled, debug messages are not logged.

Prerequisites

You must enable PIM sparse mode on the interface that is used for BGP peering. Configure PIM on all interfaces used for IP multicast. Cisco recommends configuring PIM sparse mode on all physical interfaces of PE routers connecting to the backbone. Cisco also recommends configuring PIM sparse mode on all loopback interfaces if they are used for BGP peering or if their IP address is used as an RP address for PIM.

To be able to use Auto-RP within a VRF, the interface facing the CE router must be configured for PIM sparse-dense mode.

In this example, multicast and MPLS capability have been enabled at the global and interface levels using the standard commands. A VPN named “A” has been created and multicast-enabled using the new **ip multicast-routing vrf** command. The default MDT group for a VRF is configured using the **mdt default** command. The multicast group address range for data MDT groups is configured using the **mdt data group-address-range wildcard-bits [threshold threshold-value] [list access-list]** command.

Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- **Cisco MPLS Managed Services are VPN-aware features delivered in Cisco IOS software for enabling NAT, DHCP, router redundancy, and multicast.**
- **MPLS NAT allows multiple MPLS VPNs that use the same IP addressing to share central services.**
- **MPLS DHCP relay enables a DHCP relay agent to forward a client DHCP request by adding VPN association when forwarding the packets to a shared centralized DHCP server.**
- **MPLS on-demand address pools provide the capability for automating VPN IP address prefix assignment and monitoring.**
- **HSRP and VRRP improve network availability by preventing first-hop router failure.**
- **Multicast VPNs allow ISPs to provide scalable multicast services over Layer 3 VPNs to their customers multiple sites.**

© 2004 Cisco Systems, Inc. All rights reserved. MPLS v2.1—6-31

Module Summary

This topic summarizes the key points discussed in this module.

Module Summary

Cisco.com

- **Advanced VRF features allow selective import or export of routes.**
- **Overlapping VPNs are used to provide connectivity between segments of two VPNs.**
- **Central services VPNs are used to share a common set of servers with VPNs of multiple customers.**
- **CE routers for all VPNs can be managed by service providers using a separate network management VPN.**
- **With MPLS managed services, ISPs can provide additional centralized services that are integrated with existing VPN service to customers.**

© 2004 Cisco Systems, Inc. All rights reserved.MPLS v2.1—6-5

Market forces continually drive service providers to provide more complex centralized services for their customers. These services, such as advanced VRF import and export features, overlapping VPNs, and central services VPNs, help to meet service requirements and provide VPN solutions and topologies.

References

For additional information, refer to these resources:

- Access Cisco.com for additional information about VPNs.
- NAT Integration with MPLS
VPNs: http://www.cisco.com/en/US/partner/products/sw/iosswrel/ps1839/products_feature_guide09186a00801145f5.html#22289.
- DHCP Relay—MPLS VPN Support:
http://www.cisco.com/en/US/partner/products/sw/iosswrel/ps1839/products_feature_guide09186a0080087d3c.html.
- Multicast VPN Design Guide
http://www.cisco.com/en/US/partner/tech/tk828/tk363/tech_digest09186a00801a64a3.html.
- Multicast VPN: IP Multicast Support for MPLS VPNs
http://www.cisco.com/en/US/partner/products/sw/iosswrel/ps1829/products_feature_guide09186a00801039b0.html.

Module Self-Check

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

Q1) Why do you need a selective VRF import command? (Source: Using Advanced VRF Import and Export Features)

Q2) How does the import route map affect the VRF import process? (Source: Using Advanced VRF Import and Export Features)

Q3) Why do you need a selective VRF export command? (Source: Using Advanced VRF Import and Export Features)

Q4) How does the export route map affect the VRF export process? (Source: Using Advanced VRF Import and Export Features)

Q5) Which BGP attributes can be set with an export route map? (Source: Using Advanced VRF Import and Export Features)

Q6) Who are the typical users of overlapping VPNs? (Source: Introducing Overlapping VPNs)

Q7) What are the connectivity requirements for overlapping VPNs? (Source: Introducing Overlapping VPNs)

Q8) How many VRFs do you need at most to implement three partially overlapping VPNs? How many route distinguishers? How many route targets? (Source: Introducing Overlapping VPNs)

Q9) What are the typical usages for a central services VPN topology? (Source: Introducing Central Services VPNs)

Q10) What is the connectivity model for a central services VPN topology? (Source: Introducing Central Services VPNs)

Q11) What command syntax do you use to implement a central services VPN topology that supports two clients? (Source: Introducing Central Services VPNs)

Client PE router

Server PE router

_____	_____
_____	_____
_____	_____
_____	_____
_____	_____
_____	_____

Client PE router

Server PE router

_____	_____
_____	_____
_____	_____
_____	_____

Q12) How many route distinguishers do you need for a central services VPN solution with 50 client sites and 3 server sites? How many route targets? (Source: Introducing Central Services VPNs)

route distinguishers = _____ route targets = _____

Q13) How do you combine a central services VPN topology with a simple VPN topology? (Source: Introducing Central Services VPNs)

Q14) Why do you need the managed CE routers service? (Source: Introducing Managed CE Routers Service)

Q15) What is the main difference between the managed CE routers service and the typical central services VPN topology? (Source: Introducing Managed CE Routers Service)

Q16) What syntax would you use for an export statement that limits the export to the loopback address of 192.168.10.1? (Source: Introducing Managed CE Routers Service)

Q17) What steps do you need to take to enable a VPN-aware NAT service on an existing MPLS network? (Source: Introducing MPLS Managed Services)

Q18) When you are implementing a VPN-aware DHCP relay service, which command do you use to configure the DHCP server address? (Source: Introducing MPLS Managed Services)

Q19) When you are implementing a VPN-aware DHCP relay service, how does the DHCP server know which VPN that a request comes from? (Source: Introducing MPLS Managed Services)

Q20) When you are implementing a VPN-aware ODAP service, from where does the DHCP server get its initial address pool? (Source: Introducing MPLS Managed Services)

- A) The DHCP server requests its initial address pool from the ODAP server.
- B) The DHCP server requests its initial address pool from the DHCP relay agent.
- C) The DHCP server requests its initial address pool from its local address pool via the **ip nat pool** command.
- D) The DHCP server requests its initial address pool from its local address pool via the **ip dhcp pool** command.

Q21) How is a VPN-aware HSRP or VRRP service implemented? (Source: Introducing MPLS Managed Services)

Q22) Which type of traffic flows over the default MDT? (Source: Introducing MPLS Managed Services)

Q23) What is a data MDT? (Source: Introducing MPLS Managed Services)

- Q15) The VRF and RD design is similar to that of a central services VPN. The managed CE routers service combines a service VPN and simple VPN topology like the central services VPN. However, the route export statement uses an access list to limit the exported addresses to the loopback address of the managed routers.
- Q16) ip vrf VPN_A
export route-map NMS
route-map NMS
match ip access-list 10
set extcommunity rt 123:100 additive
access-list 10 permit 192.160.10.1 0.0.0.255
- Q17) Create the NAT address pool (using the **ip nat pool** command).
Assign the address pool to the VRF (using the **ip nat inside source** command).
Enable NAT on the interfaces (using the **ip nat outside** and **ip nat inside** commands).
Create a static address to the next hop (using the **ip route vrf** command).
- Q18) **ip helper-address vrf**
- Q19) The DHCP relay agent inserts the VPN ID into the DHCP request.
- Q20) A
- Q21) by enabling a VRF on the interface and then configuring the HSRP or VRRP service using the standard commands
- Q22) Default MDT group is used for control traffic and flooding channel for dense-mode and low-bandwidth groups
- Q23) an MDT group created on demand for mVPN (S,G) pairs, usually high-bandwidth traffic

Internet Access from an MPLS VPN

Overview

Integrating Internet access with a Multiprotocol Label Switching (MPLS) Virtual Private Network (VPN) solution is one of the most common service provider business requirements. This module provides a good understanding of underlying design issues, several potential design scenarios, and some sample configurations. Various topologies and implementation methods are discussed, along with ways to separate Internet access from VPN services.

Module Objectives

Upon completing this module, you will be able to describe the various Internet access implementations available. This ability includes being able to meet these objectives:

- Describe the characteristics of MPLS VPN Internet topologies
- Describe VPN Internet access implementation methods
- Describe methods to separate Internet access from VPN services
- Describe the characteristics of implementing Internet access solutions in which Internet access is provided through a separate VPN

Introducing VPN Internet Access Topologies

Overview

This lesson identifies the characteristics of the different modules that are used to combine Internet access with VPN services.

This lesson is crucial for learners planning to enhance their usage of network resources using MPLS VPNs.

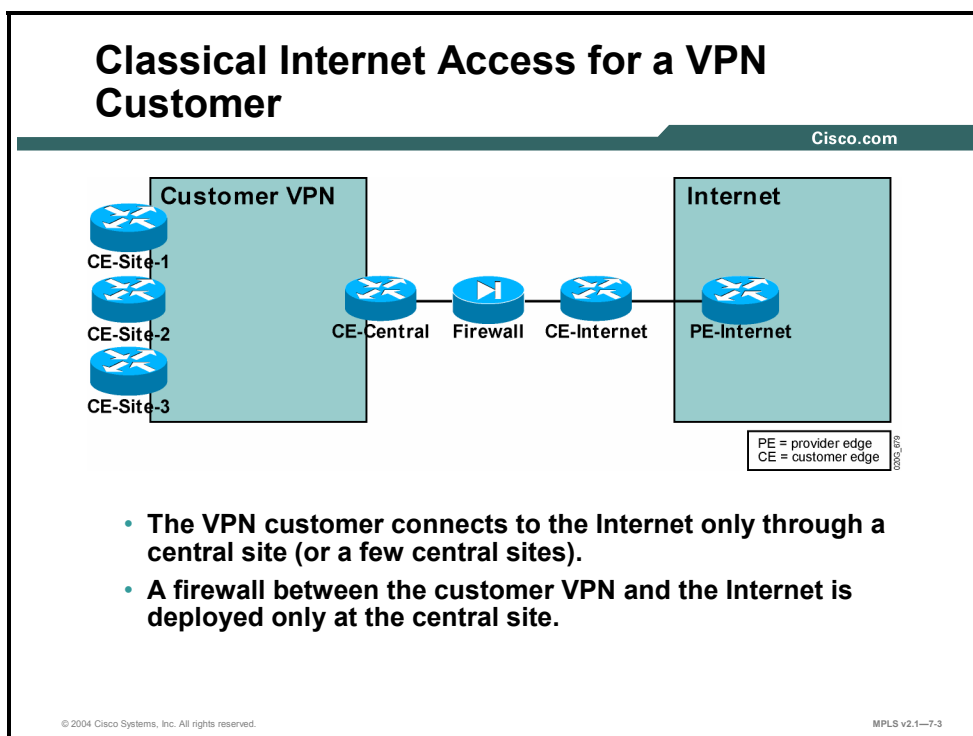
Objectives

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

- Describe the features of classical Internet access for a VPN customer
- Describe the ways in which customers may access the Internet from their own sites
- Describe the characteristics of a central firewall service
- Describe the characteristics of a wholesale Internet access service

What Is Classical Internet Access for a VPN Customer?


This topic describes the features of classical Internet access for a VPN customer.



Classical Internet access is implemented through a (usually central) firewall that connects the customer network to the Internet in a secure fashion. The private network of the customer (or VPN if the customer is using a VPN service) and the Internet are connected only through the firewall.

Classical Internet Access Addressing

Cisco.com



- The customer can use private address space.
- The firewall provides NAT between the private address space and the small portion of public address space assigned to the customer.

© 2004 Cisco Systems, Inc. All rights reserved.


MPLS v2.1—7-4

The addressing requirements for a classical Internet connection (as described here) are very simple:

- The customer is assigned a small block of public address space used by the firewall.
- The customer typically uses private addresses inside the customer network.
- The firewall performs Network Address Translation (NAT) between the private addresses of the customer and the public addresses assigned to the customer by the Internet service provider (ISP). Alternatively, the firewall might perform an application-level proxy function that isolates private and public IP addresses.

Classical Internet Access for a VPN Customer

Cisco.com



Benefits:

- The classical design is a simple, well-known setup.
- Only a single point needs to be secured.

Drawback:

- All Internet traffic from all sites goes across the central site.

© 2004 Cisco Systems, Inc. All rights reserved.


MPLS v2.1—7-5

There are a number of benefits associated with the classical design, including the following:

- It is a well-known setup used worldwide for Internet connectivity from a corporate network. Access to expertise needed to implement such a setup is thus simple and straightforward.
- There is only one interconnection point between the secure customer network and the Internet. Security of Internet access needs to be managed only at this central point.

The major drawback of this design is the traffic flow—all traffic from the customer network to the Internet has to pass through the central firewall. While this might not be a drawback for smaller customers, it can be a severe limitation for large organizations with many users, especially when the users are geographically separated.

Internet Traffic Flow in an MPLS VPN Backbone



- **Internet traffic flow becomes a more serious issue in combined VPN and Internet backbones.**
- **Some customers would like to optimize traffic flow and gain access to the Internet from every site.**

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—7-6

The traffic flow issue becomes even more pronounced when the customer VPN (based on, for example, MPLS VPN services) and the Internet traffic share the same service provider backbone. In this case, the traffic from a customer site may have to traverse the service provider backbone as VPN traffic and then return into the same backbone by the corporate firewall, ending up at a server very close to the original site.

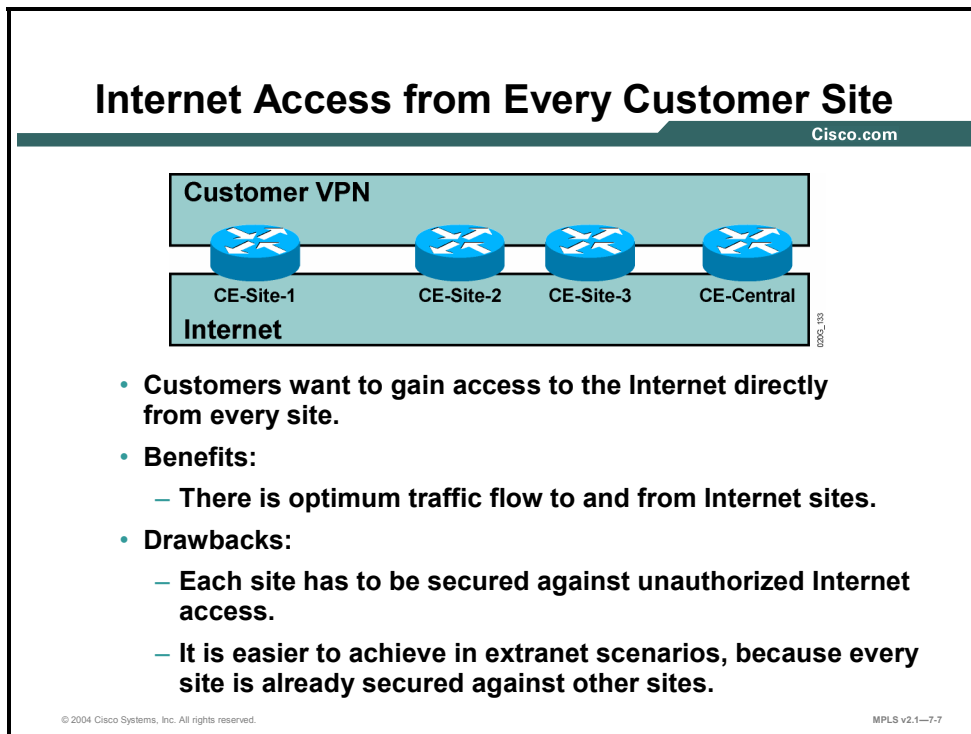
Based on this analysis, the drawbacks of the central firewall design can be summarized as follows:

- The link between the central site and the provider backbone has to be overdimensioned, because it has to transport all of the customer Internet traffic.
- The provider backbone is overutilized, because the same traffic crosses the backbone twice, first as VPN traffic and then as Internet traffic (or vice versa).
- Response times and quality of service (QoS) may suffer. This is because the traffic between the customer site and an Internet destination always has to cross the central firewall, even when the Internet destination is very close to the customer site.

These drawbacks have prompted some large users and service providers to consider alternate designs in which every customer site can originate and receive Internet traffic directly.

What Are the Methods to Access the Internet from Every Customer Site?


This topic describes ways in which customers may access the Internet from their own sites.



To bypass the limitations of Internet access through a central firewall, some customers are turning toward designs in which each customer site has its own independent Internet access. While this design clearly solves all traffic flow issues, the associated drawback is higher exposure—each site has to be individually secured against unauthorized Internet access. This design is applicable primarily for larger sites (concentrating traffic from nearby smaller sites) or for extranet VPNs, in which each site is already secured against the other sites participating in the extranet VPN.

Internet Access from Every Site: Addressing

Cisco.com



Two addressing options:

- **Every CE router performs NAT functionality—a small part of the public address space has to be assigned to each CE router.**
- **The customer uses only public IP addresses in the private network (not realistic for many customers).**

© 2004 Cisco Systems, Inc. All rights reserved.


MPLS v2.1—7-8

To gain Internet access from every site, each site requires at least some public IP addresses. The following two methods can be used to achieve this goal:

- A small part of public address space can be assigned to each customer site. NAT between the private IP addresses and the public IP addresses needs to be performed at each site.
- If the customer is already using public IP addresses in the VPN, NAT functionality is not needed. Unfortunately, this option is open only to those customers that own large address blocks of public IP addresses.

Internet Access from Every Site: MPLS VPN Backbone

Cisco.com



- Internet and VPN traffic are flowing over the provider edge-customer edge link—additional security is needed on CE routers.
- The traffic flow between an individual site and Internet destinations is always optimal.

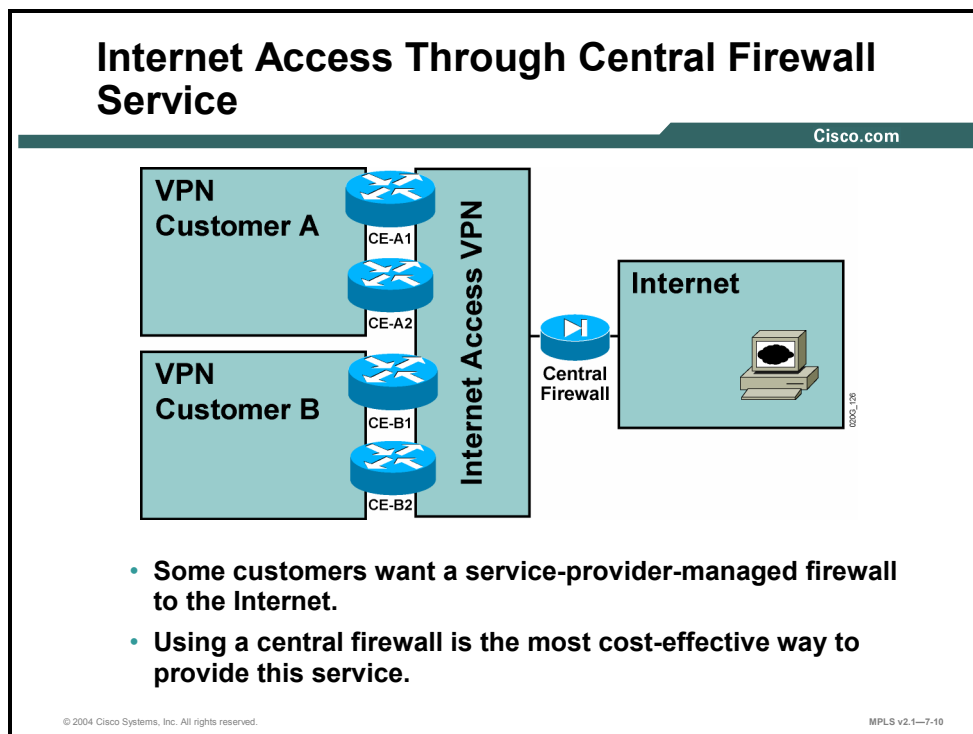
© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—7-9

To achieve Internet access from every customer site, each customer edge (CE) router must forward VPN traffic toward other customer sites and forward Internet traffic toward Internet destinations. The two traffic types are usually sent over the same physical link to minimize costs. Switched WAN encapsulation (Frame Relay or ATM) can be used to separate the VPN and Internet traffic onto different virtual circuits, or the traffic can share the same logical link as well, resulting in reduced security. On the other hand, the weaker (and less complex) security of this design is offset by optimal traffic flow between every site and Internet destinations.

What Is a Central Firewall Service?

This topic describes the characteristics of a central firewall service.

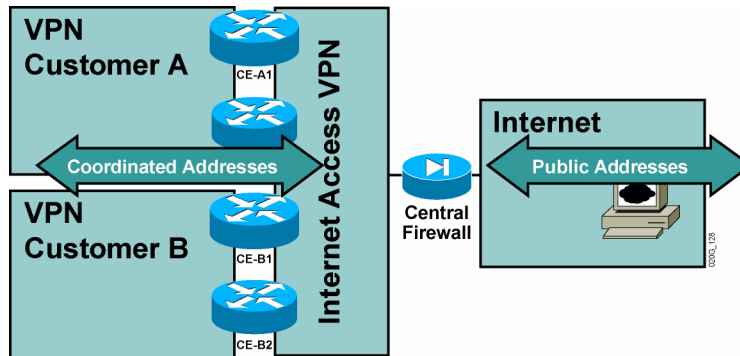


For customers who do not want the complexity of managing their own firewall, a managed firewall service offered by the service provider is a welcome relief. These customers typically want the service provider to take care of the security issues of their connection to the Internet.

The service provider can implement the managed firewall service by deploying a dedicated firewall at each customer site or (for a more cost-effective approach) by using a central firewall that provides secure Internet access to all customers.

Central Firewall Service Addressing

Cisco.com



- All customers have to use coordinated addresses, which can also be private.
- The central firewall provides NAT for all customers.

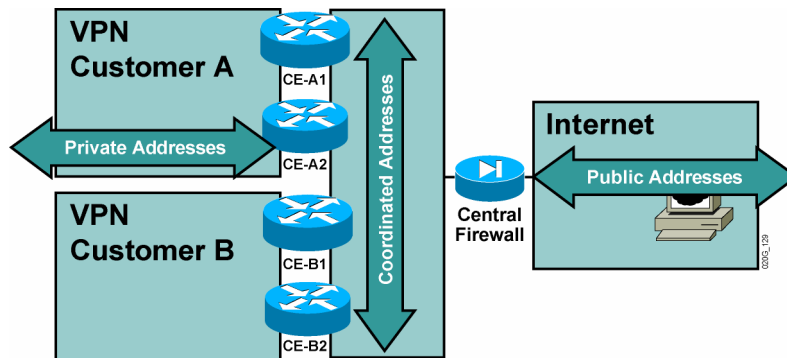
© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—7-11

The central firewall, hosted by the service provider, has to use public addresses for the Internet. Private addresses can be used between the central firewall and the individual customers. However, these addresses need to be coordinated between the service provider and the customers to prevent routing conflicts and overlapping addresses visible to the central firewall. Customers using a central firewall service are thus limited to IP addresses assigned to them by the service provider, much in the same way that Internet customers are limited to the public IP addresses assigned by their ISP.

Central Firewall Service Addressing (Cont.)

Cisco.com



- Each customer can use private address space if the CE routers provide address translation between private and coordinated address space.

© 2004 Cisco Systems, Inc. All rights reserved.


MPLS v2.1—7-12

Customers of a central firewall service who still want to retain their own private addresses inside their network can use NAT on the CE routers, connecting their private network to the transit network that links customer sites to the central firewall.

Note Service providers usually use private IP addresses as the address space between the central firewall and the customers. There is always a potential for overlapping addresses between the coordinated address space and the address space of an individual customer. Therefore, the CE device providing NAT functionality has to support address translation between overlapping sets of IP addresses.

Central Firewall Service: Traffic Flow

Cisco.com



- Traffic can flow from customer sites to the Internet and back; customer sites are protected by a central firewall.
- Traffic between sites of one customer should flow inside the customer VPN.
- Traffic between customers is not allowed; a security breach could occur.

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—7-13


The traffic flow between sites participating in a central firewall service is limited by the security requirements of the service, as described here:

- Traffic between the customer sites and the Internet must flow freely, restricted only by the security functions of the central firewall.
- Traffic between sites of an individual customer should never flow across the VPN that links the customer sites with the central firewall. This traffic must flow inside the customer VPN.
- Traffic between customers using the central firewall is not allowed, because the individual customers are not protected from outside access. (This is the task of the service provider, handled by the central firewall.) Intercustomer traffic could lead to potential security problems.

Note The restrictions on intercustomer traffic prevent customers from deploying publicly accessible servers in their networks—because these servers would not be available to other customers of the same service.

What Is Wholesale Internet Access?

This topic describes the characteristics of a wholesale Internet access service.

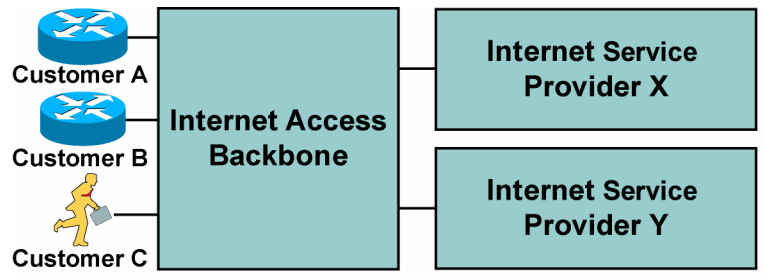


Similar to a wholesale dial service (where an ISP uses the modem pools of other service providers) is the wholesale Internet access service, where an ISP uses the IP transport infrastructure of another service provider to reach the end users. The business model of this service varies—the end users might be customers of the service provider that owns the transport backbone (for example, a cable operator), who offers Internet access through a large set of ISPs as a value-added service. Alternatively, the service provider owning the Internet access backbone might act as a true wholesaler, selling transport infrastructure to Internet service providers, who then charge end users for the whole package.

When a service provider owns the backbone and provides Internet access to customers, that service provider usually wants to offer a wide range of upstream ISPs to choose from—to satisfy various customer connectivity and reliability requirements. The selection of upstream ISPs and the corresponding configuration process should therefore be as easy as possible.

Wholesale Internet Access Addressing

Cisco.com



- **Customers get address space from the ISP that they connect to.**
- **When using dynamic addresses, the wholesale Internet access provider has to use a different address pool for every upstream service provider.**

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—7-15

Regardless of the business model used in the wholesale Internet access service, the addressing requirements are always the same: the upstream ISP allocates a portion of its address space to the end users connected to the Internet access backbone. The wholesale Internet access provider consequently has to use a different address pool for every upstream ISP.

Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- **Classical Internet access with VPN is implemented through a central firewall that connects the customer network to the Internet in a secure fashion.**
- **Customers can gain Internet access directly from the firewall of every site or collectively from a centralized ISP managed firewall service.**
- **ISP choice of centralized Internet service affects the design of unique IP addressing and location of the NAT service.**
- **With a wholesale Internet access service, an ISP uses the IP transport infrastructure of another service provider to reach end users.**

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—7-16

Introducing VPN Internet Access Implementation Methods

Overview

This lesson identifies different design models for combining Internet access with VPN services. The lesson lists the benefits and drawbacks of these models, and then explains the implications of their use.

This lesson is crucial for learners planning to enhance their usage of network resources using MPLS VPNs.

Objectives

Upon completing this lesson, you will be able to describe VPN Internet access implementation methods. This ability includes being able to meet these objectives:

- Identify the major design models for combining Internet access with VPN services
- Describe the implementation options of Internet access through global routing
- Describe the characteristics of separate interfaces and subinterfaces to provide Internet access
- Describe the benefits and drawbacks of Internet access in VPNs

Major Design Models

This topic identifies two major design models for combining Internet access with VPN services.

Major Design Models

Cisco.com

Two major design models:

- **Separating Internet access from VPN services**
- **Internet access as a separate VPN**

© 2004 Cisco Systems, Inc. All rights reserved. MPLS v2.1—7-3

Network designers who want to offer Internet access and MPLS VPN services in the same MPLS backbone can choose between these two major design models:

- Internet routing that is implemented through global routing on the provider edge (PE) routers
- Internet routing that is implemented as yet another VPN in the ISP network

Internet Access Through Global Routing

This topic describes the implementation options of Internet access through global routing.

Internet Access Through Global Routing

Cisco.com

Implementation option:

- **Internet access is implemented via separate interfaces that are not placed in any VRF (traditional Internet access setup).**

© 2004 Cisco Systems, Inc. All rights reserved.MPLS v2.1—7-4

Implementing Internet access through global routing is identical to building an IP backbone offering Internet services. IP version 4 (IPv4) Border Gateway Protocol (BGP) is deployed between the PE routers to exchange Internet routes, and the global routing table on the PE routers is used to forward the traffic toward Internet destinations.

VPN customers can reach the global routing table by this method: The customers can use a separate logical link for Internet access. This method is equivalent to traditional VPN and Internet access.

Internet Access Through Separate Interfaces or Subinterfaces

This topic describes the characteristics of separate interfaces or subinterfaces to provide Internet access.

Internet Access Through Separate Interfaces or Subinterfaces

Cisco.com

Benefits:

- **Well-known setup; equivalent to classical Internet service**
- **Easy to implement; offers a wide range of design options**

Drawback:

- **Requires separate physical links or WAN encapsulation that supports subinterfaces**

© 2004 Cisco Systems, Inc. All rights reserved. MPLS v2.1—7-5

Internet access through separate logical links is easy to set up, because 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 provider edge-customer edge (PE-CE) connectivity. Separation of Internet and VPN connectivity requires either two separate physical links or a single physical link with WAN encapsulation that supports subinterfaces (for example, Frame Relay).

Note Some customers might be reluctant to change their encapsulation type to Frame Relay, because the IP QoS mechanisms on Frame Relay differ from those provided on PPP links.

Internet Access in VPNs

This topic describes the benefits and drawbacks of having Internet access in VPNs.

Internet Access in VPNs

Cisco.com

Benefits:

- **The provider backbone is isolated from the Internet; increased security is realized.**

Drawbacks:

- **All Internet routes are carried as VPN routes; full Internet routing cannot be implemented because of scalability problems.**

© 2004 Cisco Systems, Inc. All rights reserved.MPLS v2.1—7-6

The major benefit of implementing Internet access as a separate VPN is increased isolation between the provider backbone and the Internet—which results in increased security. The flexibility of MPLS VPN topologies also provides for some innovative design options that allow the service providers to offer services that were simply not possible to implement with pure IP routing.

The obvious drawback of running the Internet as a VPN in the MPLS VPN architecture is the scalability of such a solution. An Internet VPN simply cannot carry full Internet routing because of the scalability problems associated with carrying close to 100,000 routes inside a single VPN.

Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- **There are two major design models for combining Internet access with VPN services:**
 - **Separating Internet access from VPN services**
 - **Internet services as a separate VPN**
- **Providing Internet access using global routing is implemented using separate interfaces that are not placed in any VRF.**
- **Internet access through separate interfaces or subinterfaces is a well-known setup and is easy to implement.**
- **Internet access via a separate VPN in the ISP network allows for service provider separation of backbone and Internet traffic.**

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—7-7

Separating Internet Access from VPN Services

Overview

There are times when cost factors prohibit separate physical links for VPN and Internet traffic. This lesson describes the characteristics of an Internet access service in which the Internet access is totally separate from MPLS VPN services. It is important to understand why you would choose to separate Internet access from VPN services so that you can implement the correct solution.

This lesson identifies the PE-CE requirements for separating Internet access from VPN services and how to implement the solution in an MPLS VPN network. This lesson is crucial for learners planning to enhance their usage of network resources using MPLS VPNs.

Objectives

Upon completing this lesson, you will be able to describe the methods to separate Internet access from VPN services. This ability includes being able to meet these objectives:

- Describe how customer Internet access is implemented over different interfaces than VPN services
- Describe how separate subinterfaces are implemented
- Describe the features of classical Internet access for a VPN customer
- Describe how Internet access is obtained from every customer site
- Identify the benefits and limitations of separate Internet access

Internet Access Separated from VPNs

This topic describes how customer Internet access is implemented over different interfaces than VPN services.

Designing Internet Access Separated from VPNs

Cisco.com

Customer Internet access is implemented over different interfaces than VPN access:

- Represents the traditional Internet access implementation model
- Requires separate physical links or separate subinterfaces
- Maximum design flexibility; Internet access totally independent from MPLS VPNs

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—7-3

Internet access can always be implemented with the traditional implementation model—with two links between the customer site or sites and the service provider network: a VPN link and an Internet link. The two links can be implemented with one physical link if you use a Layer 2 encapsulation that supports subinterfaces (Frame Relay, ATM, or a VLAN).

The traditional Internet access implementation model provides maximum design flexibility, because the Internet access is completely separated from the MPLS VPN services. Nevertheless, the limitations of traditional IP routing prevent this implementation method from being used for innovative Internet access solutions such as wholesale Internet access.

Implementing Separate Subinterfaces

This topic describes how separate subinterfaces are implemented.

Implementing Separate Subinterfaces

Cisco.com

- **Separate physical links for VPN and Internet traffic are sometimes not acceptable because of high cost.**
- **Subinterfaces can be used over WAN links using Frame Relay or ATM encapsulation (including xDSL).**
- **A tunnel interface could be used; however, there are these problems:**
 - **Tunnels are not VRF-aware: VPN traffic must run over a global tunnel.**
 - **This setup could lead to security leaks because global packets could end up in VPN space.**

© 2004 Cisco Systems, Inc. All rights reserved.MPLS v2.1—7-4


In situations where the cost factor prohibits separate physical links for VPN and Internet traffic, subinterfaces can be used to create two logical links over a single physical link. Subinterfaces can be configured only on WAN links using Frame Relay or ATM encapsulation (including xDSL) and on LAN links using any VLAN encapsulation—Inter-Switch Link Protocol (ISL) or 802.1q. For all other encapsulation types, a tunnel interface can be used between the CE router and the PE router.

However, the use of tunnel interfaces is strongly discouraged for the following security reasons:

- A tunnel interface on a PE router is not virtual routing and forwarding (VRF)-aware. The endpoints of the tunnel have to be in a global routing table—the VPN traffic must be tunneled across an Internet interface.
- It is also very easy to spoof generic routing encapsulation (GRE) tunnels (if the tunnel key is configured, and the key is known). An intruder from the Internet could easily generate traffic that could appear as if it were coming over the GRE tunnel from the CE router and would therefore be forwarded into the customer VPN.

Example of Internet Access Through a Dedicated Subinterface

Cisco.com




Example: Internet Access Through a Dedicated Subinterface

The figure illustrates the configuration needed to implement Internet access through a dedicated Frame Relay interface. The following configuration steps are performed:

- The customer virtual routing and forwarding instance (VRF) VPN-A is created.
- Frame Relay encapsulation is configured on the PE-CE link (Serial0).
- The VPN subinterface (Serial0.1) is created and associated with data-link connection identifier (DLCI) 101.
- The Internet subinterface (Serial0.2) is created and associated with DLCI 102.
- The CE router is configured as a BGP neighbor in both the global BGP process and inside the VPN in the VRF VPN-A.

Note The allowas-in feature needs to be configured on the PE router if the customer is propagating individual site routes to the Internet through BGP.

Internet Access Through a Dedicated Subinterface—Traffic Flow

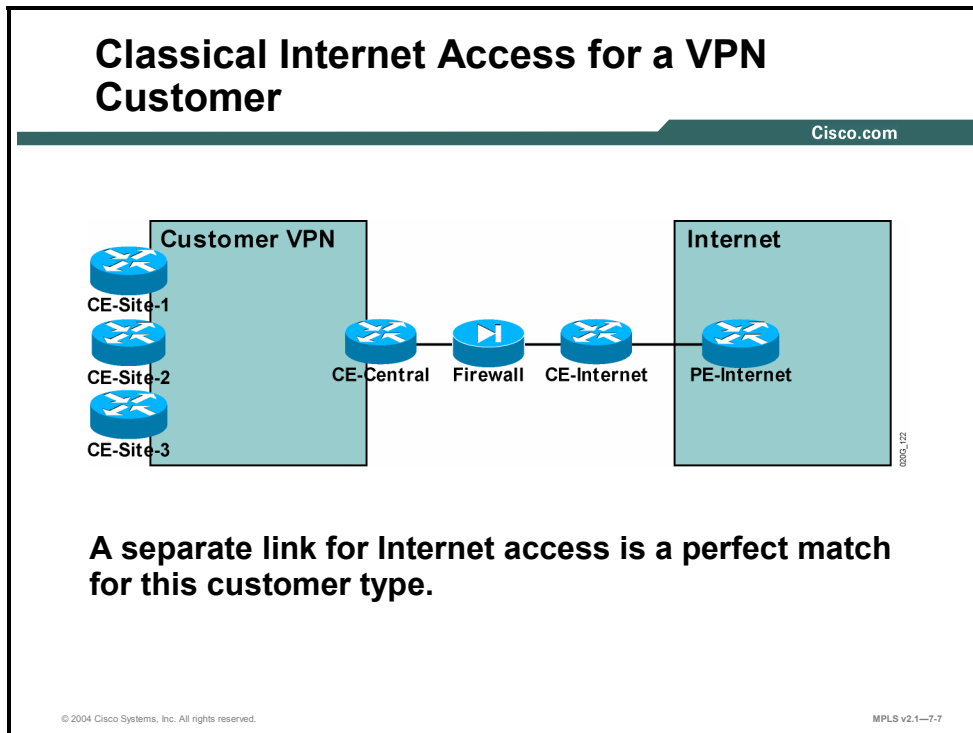


Example: Internet Access Through a Dedicated Subinterface—Traffic Flow

The Internet traffic flow in this setup is identical to the traditional Internet traffic flow—when a packet is received from the CE router through the Internet subinterface, a lookup is performed in the global Forwarding Information Base (FIB) on the PE router and the packet is forwarded toward the BGP next hop.

Classical Internet Access for a VPN Customer

This topic describes the features of classical Internet access for a VPN customer.



The classical Internet access setup for a VPN customer is based on a separated Internet access design model. This design model is thus a perfect match for customers looking for classical Internet access service.

Accessing the Internet from Every Customer Site

This topic describes how Internet access is obtained from every customer site.

Internet Access from Every Customer Site

Cisco.com

The diagram illustrates a network topology where four Customer Edge (CE) routers are connected to an Internet cloud. The routers are labeled CE-Site-1, CE-Site-2, CE-Site-3, and CE-Central. They are arranged in a horizontal line within a box labeled 'Customer VPN'. Below this box is another box labeled 'Internet'. Each router is connected to the Internet cloud by a separate link, representing a complex setup for Internet access from every site.

- Using a separate link or links for Internet access will lead to a complex setup for this customer type.
- Every CE router needs two links (or subinterfaces) to its PE router.

© 2004 Cisco Systems, Inc. All rights reserved. MPLS v2.1—7-8

For customers needing Internet access from every site, two physical (or logical) links between every CE router and its PE router might prove to be too complex or too expensive.

Separate Internet Access

This topic describes the benefits and limitations of separate Internet access.

Benefits and Limitations of Separate Internet Access

Cisco.com

Benefits:

- Well-known model
- Supports all customer requirements
- Allows all Internet services implementation, including a BGP session with the customer

Drawbacks:

- This design model requires separate physical link or specific WAN encapsulation.
- PE routers must be able to perform Internet routing (and potentially carry full Internet routing).
- Wholesale Internet access or central firewall service cannot be implemented with this model.

© 2004 Cisco Systems, Inc. All rights reserved. MPLS v2.1—7-9

The benefits of a separate Internet access design model are as follows:

- It is a well-known and widely understood model.
- It supports all customer requirements, including multihomed customer connectivity with full Internet routing.
- In addition, this model allows all Internet services implementation, including BGP sessions with customers.

The drawbacks of this model are as follows:

- It requires two dedicated physical links between the PE and the CE router or specific WAN or LAN encapsulations that might not be suitable for all customers.
- The PE routers must be able to perform hop-by-hop Internet routing and use either the default route to reach the Internet or carry the full Internet routing table.
- Advanced Internet access services (central managed firewall service or wholesale Internet access service) cannot be realized with this model at all.

Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- **Customer Internet access is implemented over different interfaces than VPN access.**
- **Subinterfaces can be used over WAN links using Frame Relay or ATM encapsulation.**
- **Classical Internet access setup for a VPN customer is based on a separated Internet access design model.**
- **For customers needing Internet access from every site, two physical (or logical) links between every CE router and its PE router might prove to be too complex or too expensive.**
- **Separate Internet access uses a well-known model that supports all customer requirements, while requiring PE routers to perform Internet routing and use separate physical or logical links.**

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—7-10

Implementing Internet Access as a Separate VPN

Overview

This lesson describes the characteristics of Internet access solutions in which the Internet access is provided as a separate VPN. The lesson identifies the scaling issues of this design and discusses how to implement the design in an MPLS VPN network.

This lesson is crucial for learners planning to improve their usage of network resources using MPLS VPNs.

Objectives

Upon completing this lesson, you will be able to describe the characteristics of implementing Internet access solutions in which the Internet access is provided as a separate VPN. This ability includes being able to meet these objectives:

- Describe the features of using Internet access as a separate VPN
- Describe how to implement a redundant Internet access implementation
- Describe how to implement classical Internet access for a VPN customer
- Describe how to implement Internet access from every customer site
- Describe how to implement Internet access through a central firewall service
- Describe how to implement wholesale Internet access
- Identify the benefits and limitations of running an Internet backbone in a VPN

What Is Internet Access as a Separate VPN?

This topic describes the features of using the Internet as a separate VPN.

Internet Access as a Separate VPN

Cisco.com

This design facilitates Internet access by using MPLS VPN features:

- **An Internet gateway is connected as a CE router to the MPLS VPN backbone.**
- **An Internet gateway shall not insert full Internet routing into the VPN; only the default route and the local (regional) routes can be inserted.**
- **Every customer that needs Internet access is assigned to the same VPN as the Internet gateway.**

© 2004 Cisco Systems, Inc. All rights reserved.


MPLS v2.1—7-3

MPLS VPN architecture suggests an obvious solution to Internet access for VPN customers: define the Internet as another VPN and use various MPLS VPN topologies to implement various types of Internet access. Under this design model, the Internet gateways appear as CE routers to the MPLS VPN backbone, and customer Internet access is enabled by combining an Internet VPN with a customer VPN in the VRFs of the customer (overlapping VPN topology).

The Internet VPN should not contain the full set of Internet routes, because that would make the solution completely nonscalable. The Internet gateway routers (CE routers) should announce a default route toward the PE routers. To optimize local routing, the local (or regional) Internet routes should also be inserted in the Internet VPN.

Internet Access as a Separate VPN (Cont.)

Cisco.com



- The Internet backbone is separate from the VPN backbone.
- VPN customers are connected to the Internet through a proper VPN VRF setup.

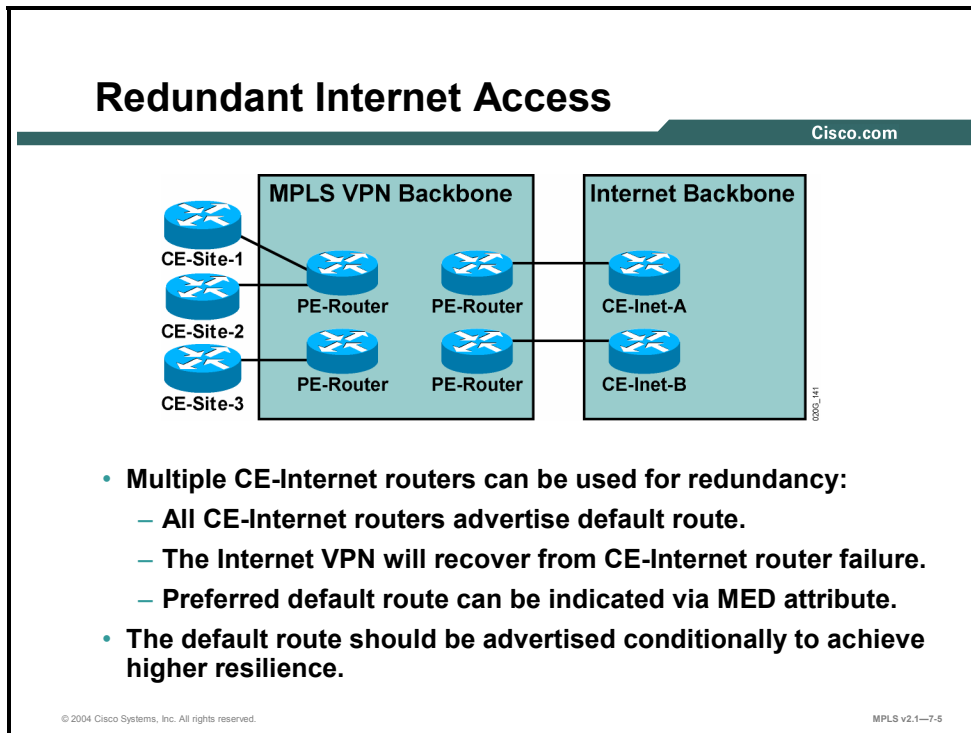
© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—7-4

When you implement Internet access as a separate VPN, the Internet backbone is kept separate from the MPLS VPN backbone—resulting in increased security for the MPLS VPN backbone (for example, Internet hosts can reach only PE routers, but not the provider routers [P routers]). The VPN customers are connected to the Internet simply through proper VRF setup.

Implementing Redundant Internet Access

This topic describes how to implement redundant Internet access.

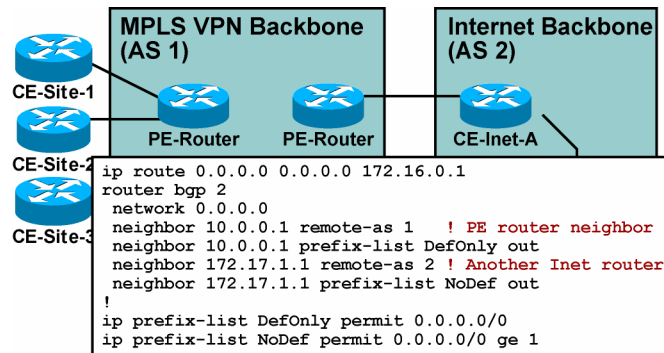


Redundant Internet access is easy to achieve when the Internet service is implemented as a VPN in the MPLS VPN backbone, as described here:

- Multiple Internet gateways (acting as CE routers) have to be connected to the MPLS VPN backbone to ensure router and link redundancy.
- All Internet gateways advertise the default route to the PE routers, resulting in routing redundancy.
- The Internet gateways also announce local Internet routes. Because these routes are announced with different BGP attributes (most notably multi-exit discriminator [MED]), the PE routers select the proper customer edge-Internet router (CE-Internet router) as the exit point toward those destinations.
- The MED attribute can also be used to indicate the preferred default route to the PE routers. In this setup, one CE-Internet router acts as a primary Internet gateway, and the other CE-Internet router acts as a backup.
- The redundancy established so far covers the path between customer sites and the CE-Internet routers. A failure in the Internet backbone might break the Internet connectivity for the customers if the CE-Internet routers announce the default route unconditionally. Conditional advertisement of the default route is therefore configured on the CE-Internet routers—the CE-Internet routers announce the default route to the PE routers only if the CE-Internet routers can reach an upstream destination.

Redundant Internet Access (Cont.)

Cisco.com



Example: CE-Inet-A should advertise the default route only if it can reach network 172.16.0.0/16 (upstream ISP core).

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—7-6

Example: Redundant Internet Access

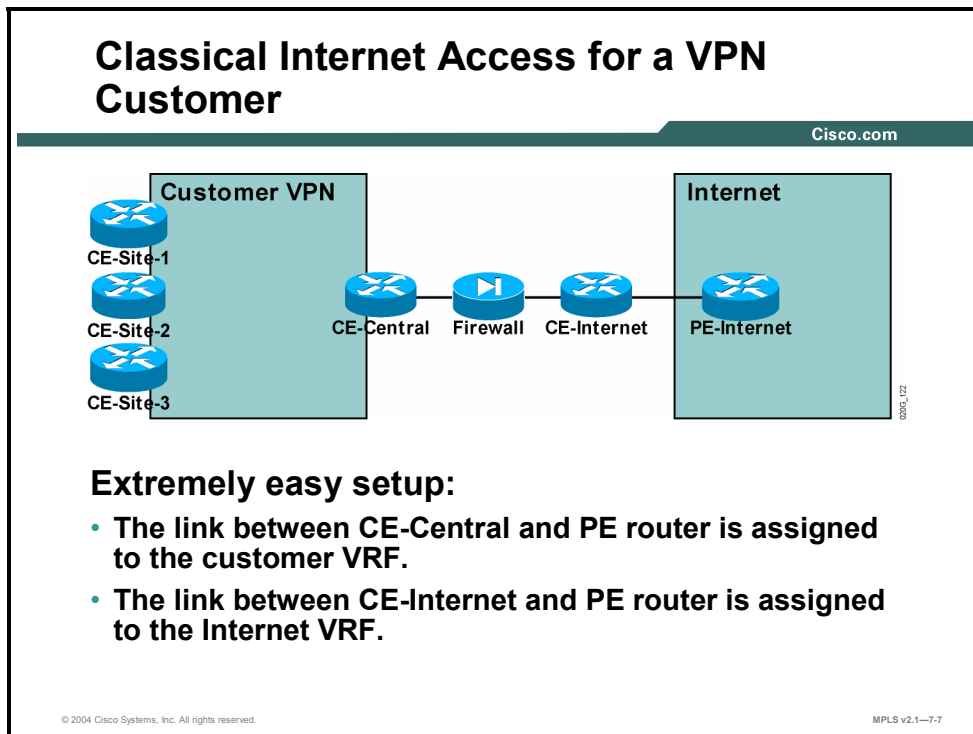
The figure shows a sample configuration of a CE-Internet router with conditional default route advertisement. Router CE-Inet-A will advertise the default route to the PE router only if it can reach the network 172.16.0.0/16.

The following steps are used to configure this functionality:

- Step 1** A static default route is configured toward a next hop in network 172.16.0.0. If the network 172.16.0.0 is not reachable, this static route will not enter the IP routing table.
- Step 2** The default route origination is configured in the BGP routing process with the network command. The default route will be originated in BGP only if it is present in the IP routing table (which, based on Step 1, means that the network 172.16.0.0/16 is reachable).
- Step 3** Prefix lists are used to filter BGP routing updates—the default route is sent only to the PE routers, not to the other Internet routers.

Implementing Classical Internet Access for a VPN Customer

This topic describes how to implement classical Internet access for a VPN customer.



The classical Internet access model can be easily implemented with the Internet configured as a VPN over the MPLS VPN backbone. The link between a PE router and the CE-Internet router is assigned to the Internet VRF, and the link between a PE router and the CE-Central router is assigned to the customer VRF. The external Border Gateway Protocol (EBGP) multihop session can be configured between the Internet gateway (CE-Internet router in the previous figure) and the CE-Internet router in this figure to give full Internet routing to the customer.

Implementing Internet Access from Every Customer Site

This topic describes how to implement Internet access from every customer site.

Internet Access from Every Customer Site

Cisco.com

The diagram illustrates a network topology where four Cisco routers are connected to two overlapping VPNs. The routers are labeled CE-Site-1, CE-Site-2, CE-Site-3, and CE-Central. They are connected to a Customer VPN (top) and an Internet VPN (bottom). The routers are connected to both VPNs, indicating an overlapping configuration.

Simple setup using overlapping VPNs:

- The customer and Internet routes are imported into the customer VRF.
- All customer routes are exported into the customer VPN.
- The public customer routes are exported into the Internet VPN.

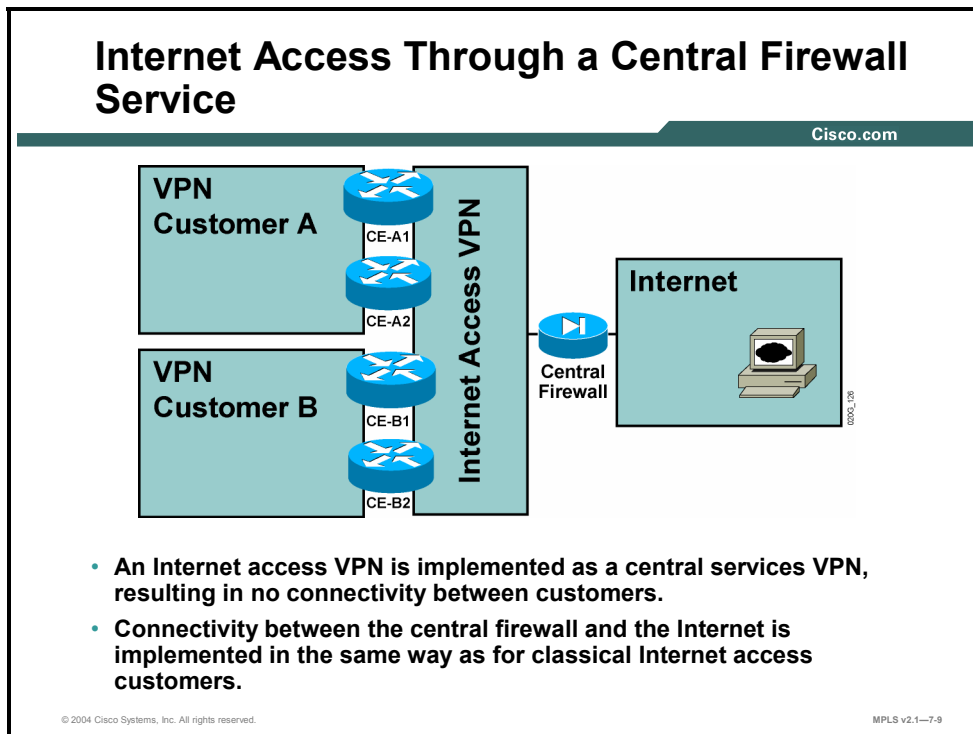
© 2004 Cisco Systems, Inc. All rights reserved. MPLS v2.1—7-8

Internet access from every customer site is best implemented with an overlapping VPN solution, as described here:

- Customer routes are marked with a customer-specific (customer) route target (RT).
- Internet routes are marked with a special (Internet) RT.
- Customer sites that need to reach the Internet are placed in a separate VRF. Customer and Internet routes are imported into this VRF, and the routes exported from this VRF are marked with customer and Internet RTs.

Implementing Internet Access Through a Central Firewall Service

This topic describes how to implement Internet access through a central firewall service.



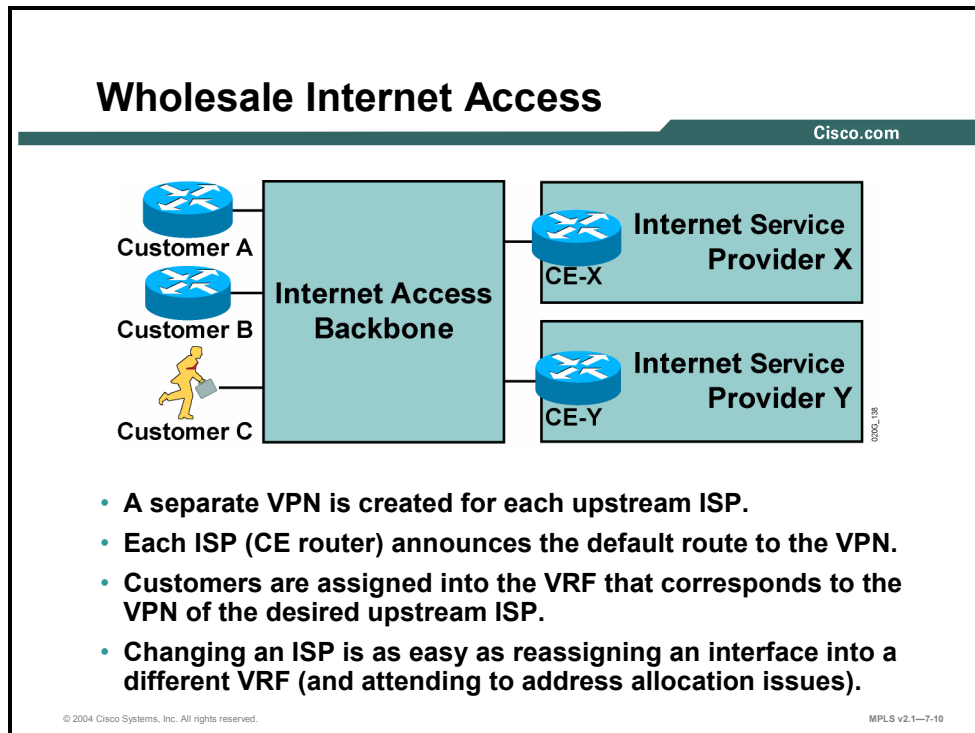
The central managed firewall service should be implemented with the central services VPN topology—with the central firewall being the server site and all customer CE routers residing in client sites. For customers with their own VPNs implemented over the same MPLS VPN backbone, the topology that overlaps customer VPN and central services VPN should be used.

The central services VPN prevents direct exchange of traffic between client sites, resulting in effective security for the customers of this service.

Connectivity between the central firewall and the Internet is implemented in the same way as Internet access for classical Internet customers. If the Internet is configured in a VPN, the public interface of the firewall is connected to an interface on a PE router, which is placed in the Internet VRF.

Implementing Wholesale Internet Access

This topic describes how to implement the wholesale Internet access model.



Wholesale Internet access is implemented by creating a separate VPN for every upstream ISP. The Internet gateway of the upstream ISP (acting as a CE router toward the MPLS VPN-based Internet access backbone) announces a default route, which is used for routing inside the VPN.

Customers are tied to upstream service providers simply by placing the PE-CE link into the VRF associated with the upstream service provider. Changing an ISP becomes as easy as reassigning the interface into a different VRF and attending to address allocation issues. For customers using access methods supporting dynamic address allocation (for example, dialup or cable), the new customer IP address is assigned automatically from the address space of the new ISP.

Running an Internet Backbone in a VPN

This topic identifies the benefits and limitations of running an Internet backbone in a VPN.

Limitations of Running an Internet Backbone in a VPN

Cisco.com

Drawbacks:

- **Full Internet routing cannot be carried in the VPN; default routes are needed that can lead to suboptimal routing.**
- **Internet backbones act as CE routers to the VPN backbone; implementing overlapping Internet + VPN backbones is tricky.**

Benefits:

- **Supports all Internet access service types**
- **Can support all customer requirements, including a BGP session with the customer, accomplished through advanced BGP setup**

© 2004 Cisco Systems, Inc. All rights reserved. MPLS v2.1—7-11

Internet access 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 has to be used, potentially resulting in suboptimal routing.

Note With future MPLS VPN extensions—called recursive VPN, or Carrier’s Carrier model—even full Internet routing will be able to be propagated across a VPN.

- The Internet backbone is positioned as a customer toward the MPLS VPN backbone. If the service provider runs Internet service and MPLS VPN service on the same set of routers, the interconnection between the two services requires special considerations.

The benefits of this design, described here, far outweigh the limitations:

- 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 the customer routes.

Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- **MPLS VPN architecture suggests an obvious solution to Internet access for VPN customers—define the Internet as yet another VPN and use various MPLS VPN topologies to implement various types of Internet access.**
- **Redundant Internet access is easy to achieve when the Internet service is implemented as a VPN in the MPLS VPN backbone.**
- **The classical Internet access model can be easily implemented with the Internet configured as a VPN over the MPLS VPN backbone.**
- **Internet access from every customer site is best implemented with an overlapping VPN solution.**
- **The central managed firewall service should be implemented with the central services VPN topology—with the central firewall being the server site and all customer CE routers residing in client sites.**
- **Wholesale Internet access is implemented by creating a separate VPN for every upstream ISP.**
- **One of the benefits of implementing Internet access as a separate VPN is that it supports all customer requirements, including full Internet routing on the customer routes.**

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.1—7-12

Module Summary

This topic summarizes the key points discussed in this module.

Module Summary

Cisco.com

- **VPN Internet topologies can be offered in a distributed fashion to every customer site or through a centralized ISP managed firewall service.**
- **Internet access can be combined with customer VPN services as a classical separate logical IPv4 interface or by the ISP adding an additional Internet VPN to enable the service.**
- **Separating Internet access from VPN services can be accomplished though use of different logical Frame Relay or ATM subinterfaces.**
- **MPLS VPN architecture defines the Internet as another VPN and uses various topologies to offer various types of Internet service.**

© 2004 Cisco Systems, Inc. All rights reserved.MPLS v2.1—7-1

There are several different models that are used to combine Internet access with VPN services. Each model has its own benefits and drawbacks. It is important to understand the implications of implementing the different models.

References

For additional information, refer to these resources:

- Access Cisco.com for additional information about MPLS Internet access.

Module Self-Check

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

- Q1) What is the major drawback of the classical Internet access model? (Source: Introducing VPN Internet Access Topologies)
- A) All of the customer traffic passes through the central firewall service.
 - B) None of the customer traffic passes through the central firewall service.
 - C) Only some of the customer traffic passes through the central firewall service.
 - D) There is no drawback.
- Q2) Which two of the following statements are NOT correct regarding Internet access from every customer site? (Choose two.) (Source: Introducing VPN Internet Access Topologies)
- A) It is easier to achieve in intranet scenarios.
 - B) It provides optimum traffic flow to and from Internet sites.
 - C) It is more difficult to achieve in extranet scenarios.
 - D) Each site has to be secured against unauthorized Internet access.
- Q3) Customers of a central firewall service who still want to retain their own private addresses inside their network can use NAT in which of the following ways? (Source: Introducing VPN Internet Access Topologies)
- A) NAT on the CE routers
 - B) NAT on the PE routers
 - C) NAT on both the CE and PE routers
 - D) NAT is not required if customers use their own private IP addresses.
- Q4) When the wholesale Internet access solution is implemented, which of the following statements is correct? (Source: Introducing VPN Internet Access Topologies)
- A) The downstream ISP allocates a portion of its address space to the end users connected to the Internet access backbone.
 - B) The upstream ISP allocates a portion of its address space to the end users connected to the Internet access backbone.
 - C) Both the upstream and downstream ISPs must allocate a portion of their address space to the end users connected to the Internet access backbone.
 - D) None of the above is correct.
- Q5) The two major design models for implementing Internet access via VPNs are using another _____ and using _____ on PE routers. (Source: Introducing VPN Internet Access Implementation Methods)
- Q6) The two major benefits of using VPNs to provide Internet access are that the _____ is isolated from the Internet and that _____ is increased. (Source: Introducing VPN Internet Access Implementation Methods)
- Q7) When you are using global routing to provide Internet access, the main implementation option is that separate interfaces not be placed in a _____. (Source: Introducing VPN Internet Access Implementation Methods)

- Q8) The main benefits of using separate interfaces or subinterfaces to provide Internet access are that they are _____ and are _____ to implement. (Source: Introducing VPN Internet Access Implementation Methods)
- Q9) The traditional Internet access implementation model provides _____ flexibility, because the Internet access is completely separated from the MPLS VPN services. (Source: Separating Internet Access from VPN Services)
- Q10) In situations where the cost factor prohibits separate physical links for VPN and Internet traffic, _____ can be used to create two logical links over a single physical link. (Source: Separating Internet Access from VPN Services)
- Q11) The _____ setup for a VPN customer is based on a separated Internet access design model. (Source: Separating Internet Access from VPN Services)
- Q12) For customers that need Internet access from every site, two physical (or logical) links between every CE router and its PE router might prove to be too _____ or too _____ to implement. (Source: Separating Internet Access from VPN Services)
- Q13) One of the drawbacks of a separate Internet access design model is that PE routers must be able to perform hop-by-hop Internet routing and use the _____ to reach the Internet or carry the full Internet routing table. (Source: Separating Internet Access from VPN Services)
- Q14) The Internet VPN should not contain the full set of _____ routes because that would make the solution completely nonscalable. (Source: Implementing Internet Access as a Separate VPN)
- Q15) All Internet gateways (CE routers) advertise the _____ route to the PE routers, which results in routing redundancy. (Source: Implementing Internet Access as a Separate VPN)
- Q16) In a classical Internet access for a customer VPN model, the link between a PE router and the CE-Internet router is assigned to the _____ VRF, and the link between a PE router and the CE-Central router is assigned to the _____ VRF. (Source: Implementing Internet Access as a Separate VPN)
- Q17) The main benefits of having Internet access from every customer site is best implemented with an _____ VPN solution. (Source: Implementing Internet Access as a Separate VPN)
- Q18) The central managed firewall service should be implemented with the _____ VPN topology. (Source: Implementing Internet Access as a Separate VPN)
- Q19) Wholesale Internet access is implemented by creating a separate VPN for every _____. (Source: Implementing Internet Access as a Separate VPN)

- Q20) One of the drawbacks of Internet access that is implemented as a separate VPN is that _____ routing cannot be carried inside a VPN. Therefore, default routing toward the Internet gateways has to be used, which can potentially result in suboptimal routing.
(Source: Implementing Internet Access as a Separate VPN)

Module Self-Check Answer Key

- Q1) A
- Q2) A, C
- Q3) A
- Q4) B
- Q5) VPN, global routing
- Q6) provider backbone, security
- Q7) VRF
- Q8) well-known, easy
- Q9) maximum design
- Q10) subinterfaces
- Q11) classical Internet access
- Q12) complex, expensive
- Q13) default route
- Q14) Internet
- Q15) default
- Q16) Internet, customer
- Q17) overlapping
- Q18) central services
- Q19) upstream ISP
- Q20) full Internet