

Implementing Cisco MPLS

Volume 2

Version 2.2

Student Guide

Text Part Number: 97-2390-01

**Corporate Headquarters**

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters

Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: 31 0 20 357 1000
Fax: 31 0 20 357 1100

Americas Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters

Cisco Systems, Inc.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

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



© 2006 Cisco Systems, Inc. All rights reserved. CCSP, CCVP, the Cisco Square Bridge logo, Follow Me Browsing, 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 Access Registrar, Aironet, 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, Cisco Unity, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, FormShare, 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, ScriptShare, SlideCast, SMARTnet, The Fastest Way to Increase Your Internet Quotient, and TransPath 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. (0601R)

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

MPLS VPN Implementation **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 VRF 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 BGP Routes—Inbound	5-11
Propagating Non-BGP Routes—Outbound	5-13
Propagating Non-BGP Routes—Inbound	5-15
Summary	5-17

Configuring VRF Tables **5-19**

Overview	5-19
Objectives	5-19
What Are the VRF Configuration Tasks?	5-20
Creating VRF Tables and Assigning RDs	5-21
ip vrf	5-21
Defaults	5-21
rd	5-22
Defaults	5-22
Specifying Export and Import RTs	5-23
route-target	5-23
Defaults	5-24
Using VPN IDs	5-25
Configuring VPN IDs	5-26
vpn id	5-26
Defaults	5-26
Assigning an Interface to a VRF Table	5-27
ip vrf forwarding	5-27
Defaults	5-27
Typical Configuration to Enable VRFs	5-28
Example: MPLS VPN Network	5-28
Summary	5-30

Configuring an MP-BGP Session Between PE Routers **5-31**

Overview	5-31
Objectives	5-31
Configuring BGP Address Families	5-32
router bgp	5-33
Defaults	5-33
address-family	5-34
Enabling BGP Neighbors	5-35
Configuring MP-BGP	5-36
Configuring MP-IBGP	5-37
neighbor remote-as	5-38
Defaults	5-38
neighbor update-source	5-38
Defaults	5-38

neighbor activate	5-39
Defaults	5-39
neighbor next-hop-self	5-40
Defaults	5-40
Configuring MP-BGP Community Propagation	5-41
neighbor send-community	5-42
Defaults	5-42
Disabling IPv4 Route Exchange	5-44
Example: Disabling IPv4 Route Exchange	5-45
Summary	5-46

Configuring Small-Scale Routing Protocols Between PE and CE Routers 5-47

Overview	5-47
Objectives	5-47
Configuring PE-CE Routing Protocols	5-48
Selecting the VRF Routing Context for BGP	5-49
address-family ipv4	5-50
Defaults	5-50
Command Modes	5-50
Configuring Per-VRF Static Routes	5-51
ip route vrf	5-51
Configuring RIP PE-CE Routing	5-53
Configuring EIGRP PE-CE Routing	5-56
Configuring SOO for EIGRP PE-CE Loop Prevention	5-59
set extcommunity	5-61
Defaults	5-62
ip vrf sitemap	5-62
Defaults	5-62
Summary	5-64

Monitoring MPLS VPN Operations 5-65

Overview	5-65
Objectives	5-65
Monitoring VRFs	5-66
show ip vrf	5-66
Defaults	5-66
Monitoring VRF Routing	5-70
show ip protocols vrf	5-70
show ip route vrf	5-71
show ip bgp vpnv4	5-72
show ip bgp vpnv4 vrf neighbors	5-75
Defaults	5-75
Usage Guidelines	5-75
Monitoring MP-BGP Sessions	5-76
show ip bgp neighbors	5-77
Example: Sample Output from show ip bgp neighbors Command	5-78
Monitoring an MP-BGP VPNv4 Table	5-80
show ip bgp vpnv4 vrf	5-81
Defaults	5-81
Usage Guidelines	5-81
show ip bgp vpnv4 rd <i>route-distinguisher</i>	5-82
Defaults	5-82
Usage Guidelines	5-83
Example: Configuring a Default RD for Two VRFs	5-83
Monitoring Per-VRF CEF and LFIB Structures	5-84
show ip cef vrf	5-85
Defaults	5-86
Usage Guidelines	5-86
show mpls forwarding vrf	5-87
Defaults	5-87

Usage Guidelines	5-87
Monitoring Labels Associated with VPNv4 Routes	5-88
Identifying Other MPLS VPN Monitoring Commands	5-89
Summary	5-90
Configuring OSPF as the Routing Protocol Between PE and CE Routers	5-91
Overview	5-91
Objectives	5-91
What Is the Enhanced OSPF Hierarchical Model?	5-92
Propagating OSPF Customer Routes	5-93
Implementing MPLS VPNs as an OSPF Superbackbone	5-96
Example: OSPF Superbackbone Implementation	5-101
Configuring OSPF PE-CE Routing	5-104
router ospf	5-106
Defaults	5-106
Using the OSPF Down Bit	5-108
Example: OSPF Down Bit	5-108
Example: OSPF Down Bit	5-110
Optimizing Packet Forwarding Across the MPLS VPN Backbone	5-111
Example: Optimizing of Packet Forwarding	5-111
Using the OSPF Tag Field	5-114
Example: Routing Loops Across OSPF Domains	5-114
Example: OSPF Tag Field—Routing Loop Prevention	5-117
What Is a Sham Link?	5-118
Example: Sham Link	5-118
Configuring a Sham Link	5-122
Defaults	5-123
Command Modes	5-124
Example: Sample Sham-Link Configuration	5-124
Summary	5-125
Configuring BGP as the Routing Protocol Between PE and CE Routers	5-127
Overview	5-127
Objectives	5-127
Configuring a Per-VRF BGP Routing Context	5-128
address-family ipv4	5-129
Defaults	5-129
Command Modes	5-129
Example: Configuring per-VRF BGP Routing Context	5-130
What Are the Reasons for Limiting the Number of Routes in a VRF?	5-131
Limiting the Number of Prefixes Received from a BGP Neighbor	5-132
neighbor maximum-prefix	5-132
Defaults	5-133
Limiting the Total Number of VRF Routes	5-134
maximum routes	5-135
Defaults	5-135
Example: Limiting the Total Number of VRF Routes	5-136
Identifying AS-Override Issues	5-137
neighbor as-override	5-140
Defaults	5-140
Example: AS-Override	5-141
Example: AS-Path Prepending	5-142
Identifying Allowas-in Issues	5-143
Example: Using Allowas-in to Support Customer Site Linking Two VPNs	5-146
neighbor allowas-in	5-147
Defaults	5-147
Implementing SOO for Loop Prevention	5-148
set extcommunity	5-150
Defaults	5-151
neighbor route-map	5-151

ip vrf sitemap	5-152
Defaults	5-152
Summary	5-154
Troubleshooting MPLS VPNs	5-155
Overview	5-155
Objectives	5-155
Identifying Preliminary Steps in MPLS VPN Troubleshooting	5-156
Verifying the Routing Information Flow	5-157
Validating CE-to-PE Routing Information Flow	5-158
Validating PE-to-PE Routing Information Flow	5-159
Validating PE-to-CE Routing Information Flow	5-164
Identifying the Issues When Verifying the Data Flow	5-165
Validating CEF Status	5-166
show cef interface	5-167
Usage Guidelines	5-167
Validating the End-to-End LSP	5-170
Validating the LFIB Status	5-171
Summary	5-172
Module Summary	5-173
References	5-174
Module Self-Check	5-175
Module Self-Check Answer Key	5-184
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 Central Services and Simple VPNs	6-33
Identifying the RT Requirements When Integrating Central Services and Simple VPNs	6-34
Example: Configuring VRFs in a Central Services and Simple VPN	6-36
Summary	6-37
Introducing the 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
Module Summary	6-45
References	6-45
Module Self-Check	6-46
Module Self-Check Answer Key	6-50
Internet Access and MPLS VPNs	7-1
Overview	7-1
Module Objectives	7-1
Introducing Internet Access with MPLS VPNs	7-3
Overview	7-3
Objectives	7-3
Customer Internet Connectivity Scenarios	7-4
Classical Internet Access	7-4
Multisite Internet Access	7-5
Wholesale Internet Access	7-6
Internet Design Models for Service Providers	7-7
Major Design Models	7-8
Internet Access Through Global Routing	7-9
Internet Access as a Separate VPN	7-10
Disadvantages of Providing Internet Access Through Route Leaking	7-11
Summary	7-13
Implementing Separate Internet Access and VPN Services	7-15
Overview	7-15
Objectives	7-15
Classical Internet Access for a VPN Customer	7-16
Using Separate Subinterfaces	7-17
Example: Internet Access Through Static Routes	7-18
Example: Dynamic Internet Access Through a Separate Subinterface	7-19
Example: Internet Access Through a Dedicated Subinterface—Traffic Flow	7-20
Accessing the Internet from Every Customer Site	7-21
Separate Internet Access Benefits and Limitations	7-22
Summary	7-23

Implementing Internet Access as a Separate VPN	7-25
Overview	7-25
Objectives	7-25
Internet Access as a Separate VPN	7-26
Example: Configuring the Internet Gateway in a Separate VPN	7-28
Implementing Redundant Internet Access	7-29
Implementing Classical Internet Access for a VPN Customer	7-30
Implementing Internet Access from Every Customer Site	7-32
Implementing Wholesale Internet Access	7-33
Running an Internet Backbone in a VPN	7-34
Summary	7-35
Module Summary	7-36
References	7-36
Module Self-Check	7-37
Module Self-Check Answer Key	7-40
MPLS TE Overview	8-1
Overview	8-1
Module Objectives	8-1
Introducing the TE Concept	8-3
Overview	8-3
Objectives	8-3
What Is TE?	8-4
Business Drivers for TE	8-6
Congestion Avoidance and TE	8-8
TE with a Layer 2 Overlay Model	8-9
TE with a Layer 3 Model	8-12
TE with the MPLS TE Model	8-14
Summary	8-16
References	8-16
Understanding MPLS TE Components	8-17
Overview	8-17
Objectives	8-17
Traffic Tunnels: Concepts	8-18
Traffic Tunnels: Characteristics	8-20
Traffic Tunnels: Attributes	8-21
Network Links and Link Attributes	8-23
Constraint-Based Path Computation	8-24
TE Processes	8-28
Role of RSVP in Path Setup and Trunk Admission Control	8-30
Path Setup with RSVP	8-31
Trunk Admission Control with RSVP	8-32
Forwarding Traffic to a Tunnel	8-33
Forwarding Traffic to a Tunnel: Autoroute	8-34
Summary	8-36
References	8-36
Configuring MPLS TE on Cisco IOS Platforms	8-37
Overview	8-37
Objectives	8-37
MPLS TE Configuration Road Map	8-38
Enabling Device-Level MPLS TE Support	8-39
ip cef	8-39
mpls traffic-eng tunnels (global)	8-40
Enabling MPLS TE Support in IS-IS	8-41
mpls traffic-eng	8-41
mpls traffic-eng router-id	8-42

metric-style wide	8-42
Enabling MPLS TE Support in OSPF	8-44
mpls traffic-eng area	8-44
mpls traffic-eng router-id	8-45
Enabling Basic MPLS TE on an Interface	8-47
mpls ip	8-47
mpls traffic-eng tunnels (interface)	8-47
ip rsvp bandwidth	8-48
Creating and Configuring a Traffic Tunnel	8-50
interface tunnel	8-50
ip unnumbered	8-50
tunnel destination	8-51
Configuring a Traffic Tunnel	8-52
tunnel mode mpls traffic-eng	8-53
tunnel mpls traffic-eng bandwidth	8-53
tunnel mpls traffic-eng priority	8-53
ip explicit-path	8-55
next-address	8-55
tunnel mpls traffic-eng path-option	8-55
Mapping Traffic into Tunnels with Autoroute	8-57
tunnel mpls traffic-eng autoroute announce	8-57
Summary	8-59
References	8-59
Monitoring Basic MPLS TE on Cisco IOS Platforms	8-61
Overview	8-61
Objectives	8-61
Monitoring MPLS TE Tunnels	8-62
show ip rsvp interface	8-62
show mpls traffic-eng tunnels	8-63
Monitoring MPLS TE	8-66
show mpls traffic-eng autoroute	8-66
show ip cef <i>network</i>	8-68
show ip cef vrf <i>vrf-name network</i>	8-69
Summary	8-70
References	8-70
Module Summary	8-71
References	8-71
Module Self-Check	8-72
Answer Key	8-75

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 instances 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 MPLS VPN operations

Using MPLS VPN Mechanisms of Cisco IOS Platforms

Overview

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

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

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 inbound 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 non-BGP route propagation process in an MPLS VPN implementation

What Is a VRF Table?

This topic describes the characteristics of a VRF table.

VRF Table

- **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

© 2006 Cisco Systems, Inc. All rights reserved. MPLS v2.2—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 a routing and forwarding instance that you can use for a single VPN site or for many sites connected to the same provider edge (PE) router if and only if these sites share exactly the same connectivity requirements.

Other MPLS VPN attributes 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 VPN version 4 (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), Intermediate System-to-Intermediate System (IS-IS), 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

Routing context = routing protocol run in one VRF:

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

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—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 (in 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

- **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**

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—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 Information Base (FIB) table 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 12.1T. 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, 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 next 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-BGP-A is the CE router for customer A and is associated with VRF-A (VPN-A). CE-BGP-B is the CE router for customer B and is associated with VRF-B (VPN-B).

BGP Route Propagation—Outbound (Cont.)



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


© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—5-8

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

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

BGP Route Propagation—Outbound (Cont.)



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

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—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 BGP routes received from BGP-speaking CE routers are copied into the MP-BGP table for further propagation to other PE routers. This is the export process.

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

Note There is not a separate per-VRF BGP table and global MP-BGP table in Cisco IOS software.


Propagating BGP Routes—Inbound

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



As other PE routers start originating VPNv4 routes, the MP-BGP process in the PE router receives the routes. The routes are filtered based on 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.

BGP Route Propagation—Inbound (Cont.)



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


© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—5-11

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.


Propagating Non-BGP Routes—Outbound

This topic describes the outbound non-BGP route propagation process in an MPLS VPN implementation. The example will discuss the case of RIP-speaking CE routers, but a similar process would support other non-BGP protocols.



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

Non-BGP Route Propagation—Outbound (Cont.)



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

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—5-13


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

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

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 Non-BGP Routes—Inbound

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



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.

Non-BGP Route Propagation—Inbound (Cont.)



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

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—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 that were discussed in this lesson.

Summary

- A VRF table is a routing and forwarding instance that associates additional attributes such as RD, import RT, and export RT to routing entries.
- 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.

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—5-16

Summary (Cont.)

- Outbound BGP route propagation starts with CE BGP updates. Because the protocol source is BGP, MP-BGP can directly prepend RDs and RTs to the respective inbound instances of CE BGP updates.
- Inbound BGP route propagation filters routes based on RT into respective instances of VRF.
- 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 non-BGP route propagation filters routes based on RT into respective VRF instances. Redistribution is required for route propagation with non-BGP speaking CEs.

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—5-17

Configuring VRF Tables

Overview

This lesson explains how to configure virtual routing and forwarding (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 Virtual Private Network (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
- Describe the optional use of VPN IDs
- Assign an interface to a VRF table
- Describe a typical Cisco IOS configuration that enables VRF

What Are the VRF Configuration Tasks?

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

VRF Configuration Tasks

VRF configuration tasks:

- **Create a VRF table**
- **Assign RD to the VRF**
- **Specify export and import route targets**
- **(Optional) Configure a VPN ID**
- **Assign interfaces to VRFs**

© 2006 Cisco Systems, Inc. All rights reserved. MPLS v2.2—5-3

Configuring a VRF table and starting deployment of a Multiprotocol Label Switching (MPLS) VPN service for a customer consists of these four mandatory steps:

- Create a new VRF table.
- Assign a unique route distinguisher (RD) to the VRF.

Note You must assign a unique RD to every VRF created in a provider edge (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 route targets (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

```
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.

© 2006 Cisco Systems, Inc. All rights reserved. MPLS v2.2—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 IP version 4 (IPv4) prefix to create a VPN version 4 (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. After local interfaces are bound to the VRF with the ip vrf forwarding command, the interfaces 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

```
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 12.4(T) and earlier.

© 2006 Cisco Systems, Inc. All rights reserved. MPLS v2.2—4-5

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 <i>route-target-ext-community</i> field that will be imported into the VRF
export	The value in the <i>route-target-ext-community</i> 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 <i>route-target-ext-community</i> 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.)

```
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 65173:15  
route-target export 65173:15  
route-target import 65173:15
```

© 2006 Cisco Systems, Inc. All rights reserved. MPLS v2.2—5.4

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

Using VPN IDs

This topic defines VPN identifiers (VPN IDs) and discusses how to configure them.

What Is a VPN ID?

- **A VPN identifier (VPN ID) allows you to identify VPNs by an ID number**
 - **Not used to control distribution of routing information**
 - **Not used to associate IP addresses with VPN IDs in routing updates**
 - **Is stored on the VRF structure for a VPN**
- **Has the following elements:**
 - **OUI (three-octet hex number)**
 - **A VPN index (four-octet hex number identifying VPN within the company)**
- **Configure all PE routers that belong to the same VPN with the same VPN ID.**
- **Make the VPN ID unique to the Service Provider network**

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—5-7

Multiple VPNs can be configured in a router. You can use a VRF name (a unique ASCII string) to reference a specific VPN configured in the router. Alternately, you can use a VPN ID to identify a particular VPN in the router as described in RFC 2685. The VPN ID is stored in the corresponding VRF structure for the VPN.

Note

Configuration of a VPN ID for a VPN is optional. You can still use a VRF name to identify configured VPNs in the router. The VPN name is not affected by the VPN ID configuration. These are two independent mechanisms to identify VPNs.

The MPLS VPN ID feature is not used to control the distribution of routing information or to associate IP addresses with MPLS VPN ID numbers in routing updates.

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

- An Organizational Unique Identifier (OUI), a three-octet hex number. The IEEE Registration Authority assigns OUIs to any company that manufactures components under the International Organization for Standardization (ISO)/International Electrotechnical Commission (IEC) 8802 standard. The OUI is used to generate universal LAN MAC addresses and protocol identifiers for use in local and metropolitan area network applications. For example, an OUI for Cisco Systems is 00-03-6B (hex).
- A VPN index, a four-octet hex number, which identifies the VPN within the company.

To ensure that the VPN has a consistent VPN ID, assign the same VPN ID to all the routers in the service provider network that services that VPN.

You can use several applications to manage VPNs using the VPN ID, such as Dynamic Host Configuration Protocol (DHCP) and RADIUS server.

Configuring VPN IDs

This section discusses how to configure VPN IDs.

Configuring VPN IDs

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

- This command creates a VRF routing table and a CEF forwarding table, and enters VRF configuration mode.

```
Router(config-vrf)#  
vpn id oui:vpn-index
```

- This command assigns the VPN ID to the VRF.

© 2006 Cisco Systems, Inc. All rights reserved. MPLS v2.2—5-8

vpn id

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

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

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

Syntax Description

Parameter	Description
<i>oui</i>	This parameter is an OUI. The IEEE organization assigns this identifier to companies. The OUI is restricted to three octets.
<i>vpn-index</i>	This value identifies the VPN within the company. This VPN index is restricted to four octets.

Defaults

By default, the VPN ID is not set.

Assigning an Interface to a VRF Table

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

Assigning an Interface to a VRF Table

```
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
```

© 2006 Cisco Systems, Inc. All rights reserved. MPLS v2.2—5.9

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.

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

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 Border Gateway Protocol (BGP) and Routing Information Protocol (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.)



The configuration steps that you 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 this MPLS VPN backbone, because these customers require only simple VPN connectivity. To simplify the configuration and troubleshooting process, the RTs are made equal to the RDs.
- Step 3** Assign PE-CE interfaces to individual VRFs.

Summary

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

Summary

There are four required VRF configuration tasks:

- **Create a VRF table**
 - Use the `ip vrf` command
- **Assign RD to the VRF**
 - Use the `rd` command
- **Specify export and import RTs**
 - Use the `route-target` command
- **Assign interfaces to VRFs.**
 - Use the `ip vrf forwarding` command, and reconfigure the IP address

Configuring a numeric VPN ID is optional.

- Use the `vpn id` command

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—5-12

Configuring an MP-BGP Session Between PE Routers

Overview

This lesson explains the Border Gateway Protocol (BGP) process in a Multiprotocol Label Switching (MPLS) Virtual Private Network (VPN)-enabled router, listing the configuration tasks, steps, syntax, and descriptions. The lesson also discusses BGP community propagation and provides a Multiprotocol Internal Border Gateway Protocol (MP-IBGP) configuration example.

Most of the configuration in an MPLS VPN depends on how the provider edge (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 Multiprotocol Border Gateway Protocol (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

- **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 External Border Gateway Protocol sessions.
- **Address families (routing protocol contexts) are used to configure these three tasks in the same BGP process.**

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—5.3

Independently from the MPLS VPN architecture, the PE router can use BGP IP version 4 (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:

- VPN version 4 (VPNv4) routes are propagated across an MPLS VPN backbone using MP-BGP between the PE routers.
- BGP can be used as the provider edge-customer edge (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.)

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**

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—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 virtual routing and forwarding (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 autonomous system (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, Routing Information Protocol (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 route distinguisher (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

- **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.**

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—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

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).**

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—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

Router(config)#

```
router bgp as-number
neighbor ip-address remote-as as-number
neighbor ip-address update-source interface-type
                             interface-number
```

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

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—6-7

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.)

```
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.

© 2006 Cisco Systems, Inc. All rights reserved.MPLS v2.2—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

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

```
neighbor ip-address send-community [standard | extended  
| both]
```

- **This command with the extended option is enabled by default by Cisco IOS software after the BGP neighbor has been activated for VPNv4 route exchange.**
- **The command can be used to enable propagation of standard BGP communities attached to VPNv4 prefixes.**
- **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.**

© 2006 Cisco Systems, Inc. All rights reserved.MPLS v2.2—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.)



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

```
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.

© 2006 Cisco Systems, Inc. All rights reserved.MPLS v2.2—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.)

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

```
router bgp 65173
  no bgp default ipv4-unicast
  neighbor 172.16.32.14 remote-as 65173
  neighbor 172.16.32.15 remote-as 65173
  neighbor 172.16.32.27 remote-as 65173

! Activate IPv4 route exchange
  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
```

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—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:

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

Summary

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

Summary

- **Use the address-family command to select the routing context that you want to configure.**
- **Use the router bgp command to configure the BGP routing process, and configure VRF-specific EBGP neighbors 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.

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—5-13

Summary (Cont.)

- **These commands are used to configure MP-IBGP:**
 - neighbor remote-as
 - neighbor update-source
 - neighbor activate
 - 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:**
 - no neighbor activate command
 - no bgp default ipv4-unicast command.

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—5-14

Configuring Small-Scale Routing Protocols Between PE and CE Routers

Overview

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

It is important to understand not only what you can configure between provider edge (PE) and customer edge (CE) routers when you are setting up Multiprotocol Label Switching (MPLS) Virtual Private Networks (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 on the PE router
- 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

- **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.**
- **Prior to Cisco IOS Release 12.3(4)T, the overall number of routing processes per router was limited to 32, of which only 28 were available for VRF assignment.**

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—5.3

After you configure virtual routing and forwarding instances (VRFs) and establish Multiprotocol Internal Border Gateway Protocol (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 PEvirtual-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 Border Gateway Protocol [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, Open Shortest Path First [OSPF]).

Note Prior to Cisco IOS Release 12.3(4)T, 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 was therefore limited to 28.

This restriction was removed for MPLS in Cisco IOS Release 12.3(4)T.

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

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 External Border Gateway Protocol neighbors in VRF context, not in global BGP configuration.**
- **All non-BGP per-VRF routes have to be redistributed into a per-VRF BGP context to be propagated by MP-BGP to other PE routers.**

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—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 IP version 4 (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

Router(config)#

```
ip route vrf vrf-name prefix mask [interface interface-  
number] [next-hop-address]
```

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

Sample router configuration:

```
ip route vrf Customer_ABC 10.0.0.0 255.0.0.0 serial0/0 10.250.0.2  
!  
router bgp 65173  
address-family ipv4 vrf Customer_ABC  
redistribute static
```

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—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 [interface {interface-number}] [next-hop-address] [global] [distance] [permanent] [tag tag]*
- **no ip route vrf** *vrf-name prefix mask [interface {interface-number}] [next-hop-address] [global] [distance] [permanent] [tag tag]*

Note You must specify a next-hop IP address if you are not using a point-to-point interface.

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>	(Optional) Type of network interface to use: ATM, Ethernet, loopback, Packet over SONET (POS), or null
<i>interface-number</i>	(Optional) 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

- 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.

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—5-6

Configuring RIP as the PE-CE routing protocol is even easier than configuring BGP. Start the configuration of the individual routing context with the **address-family ipv4 vrf vrf-name** command in router configuration mode. You can enter all standard RIP parameters in the per-VRF routing context. Global RIP parameters 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) not RIP version 1 (RIPv1) 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

Router(config)#

```
router rip
version 2
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 MED.**
 - **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.**

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—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



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

- **Provides EIGRP with the capability to redistribute routes through a VPN cloud**
- **Requires configuration of only the PE routers**
- **Imposes no upgrade or configuration changes to customer equipment**
- **Supports SOO capabilities to filter VPN traffic**

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—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 such as Site of Origin (SOO).

Configuring EIGRP PE-CE Routing: EIGRP Metric Propagation

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.

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—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




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 autonomous system (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 the customer routing context EIGRP with metric propagation is configured.
- The BGP routing context is configured for every VRF. Redistribution of the customer routing context 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.

Configuring SOO for EIGRP PE-CE Loop Prevention

This topic describes how to configure EIGRP MPLS VPN PE-CE SOO for an EIGRP PE-CE routing session to prevent routing loops.



In this multihomed scenario, a routing loop exists.

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

Implementing EIGRP SOO for Loop Prevention (Cont.)

- The SOO extended community can be used to prevent loops in dual-homed scenarios.
- All PE routers supporting EIGRP MPLS VPNs must support the SOO extended community.
- A unique SOO value must be configured for each VPN site. This value must be used on the PE-CE interface.
- The SOO attribute is configured through a route-map command.

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—5-13

The EIGRP MPLS VPN PE-CE SOO feature provides SOO support for EIGRP-to-BGP and BGP-to-EIGRP redistribution. The SOO extended community is a BGP extended community attribute that is used to identify routes that have originated from a site so that the readvertisement of that prefix back to the source site can be prevented. The SOO extended community uniquely identifies the site from which a PE router has learned a route. SOO support provides the capability to filter MPLS VPN traffic on a per-EIGRP site basis. SOO filtering is configured at the interface level and is used to manage MPLS VPN traffic and to prevent routing loops from occurring in complex and mixed network topologies, such as EIGRP VPN sites that contain both VPN and backdoor links.

The configuration of the SOO extended community allows MPLS VPN traffic to be filtered on a per-site basis. The SOO extended community is configured in an inbound BGP route map on the PE router, and is applied to the interface with the **ip vrf sitemap** command. The SOO extended community can be applied to all exit points at the customer site for more specific filtering, but must be configured on all interfaces of PE routers that provide VPN services to CE routers.

Implementing EIGRP SOO for Loop Prevention (Cont.)

Router(config)#

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

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

Router(config-if)#

```
ip vrf forwarding vrf-name  
ip vrf sitemap route-map-name  
ip address ip-address subnet-mask
```

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

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2--5-14

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 route target (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 these combinations: ■ <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.

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 EIGRP SOO for Loop Prevention (Cont.)



In this example, a route map SOO_Support with a specific SOO value of 101:2 was defined. The newly defined route map is then applied to vrf Customer-EIGRP-A1 that connects the EIGRP neighbor (CE-EIGRP-A2 router) to the PE-Site-Y router.

Summary

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

Summary

- **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.**
- **Use the `SOO` extended community to prevent loops in EIGRP dual-homed scenarios.**

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—5-16

Monitoring MPLS VPN Operations

Overview

This lesson presents the commands, syntax, and descriptions for monitoring virtual routing and forwarding (VRF) routing, Multiprotocol Border Gateway Protocol (MP-BGP) sessions, and Virtual Private Network (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 a Multiprotocol Label Switching (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

```
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

© 2006 Cisco Systems, Inc. All rights reserved. MPLS v2.2—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) Displays the 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

```
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#
```

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—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 route distinguisher (RD)
Interfaces	Specifies the network interfaces

Monitoring VRFs: show ip vrf detail

```
Router#show ip vrf detail
VRF SiteA2; default RD 103:30
  Interfaces:
    Serial11/0.20
    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:
    Serial11/0.100
    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
```

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2-5-5

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 route target (RT) export communities
Import	Specifies VPN RT import communities

Monitoring VRFs: show ip vrf interfaces

```
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

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—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

```
Router#  
show ip protocols vrf vrf-name
```

- Displays the routing protocols configured in a VRF

```
Router#  
show ip route vrf vrf-name
```

- Displays the VRF routing table

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

- Displays per-VRF BGP parameters

© 2006 Cisco Systems, Inc. All rights reserved. MPLS v2.2—5-7

These 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 Border Gateway Protocol (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 the 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 the 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 these keywords: bgp , egp , eigrp , hello , igrp , isis , ospf , or rip
<i>as-number</i>	(Optional) Specifies the autonomous system (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 VPN version 4 (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) Displays the 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) Displays the IP address of a network in the BGP routing table
<i>mask</i>	(Optional) Displays the 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) Displays 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

```

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 65031
  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)

```

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 ID 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

```
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 ...
```

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2-5.9

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

```
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 ...
```

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2--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
vpnvp4	Specifies VPNv4 information
all	Displays the complete VPNv4 database
vrf <i>vrf-name</i>	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

Router#

```
show ip bgp neighbors
```

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

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—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

```
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
```

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—5-12

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

```
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
```

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2--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) Displays the 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 External Border Gateway Protocol (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.

Here 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
                        Sent          Rcvd
Prefix activity:      ----          ----
```



```

Prefixes Current:          0          30 (Consumes 1440 bytes)
Prefixes Total:           0          30
Implicit Withdraw:        0          0
Explicit Withdraw:        0          0
Used as bestpath:         n/a        30
Used as multipath:        n/a        0
                                Outbound  Inbound
Local Policy Denied Prefixes: -----
Bestpath from this peer:   30          n/a
Total:                    30          0
Number of NLRIs in the update sent: max 0, min 0

For address family: VPNv4 Unicast
BGP table version 30, neighbor version 30
Index 4, Offset 0, Mask 0x10
NEXT_HOP is always this router
Community attribute sent to this neighbor
                                Sent      Rcvd
Prefix activity:          ----      ----
Prefixes Current:         9          1 (Consumes 256 bytes)
Prefixes Total:           18         1
Implicit Withdraw:        9          0
Explicit Withdraw:        0          0
Used as bestpath:         n/a        4
Used as multipath:        n/a        0
                                Outbound  Inbound
Local Policy Denied Prefixes: -----
VPN Imported prefix:       3          n/a
Bestpath from this peer:   1          n/a
Total:                     4          0
Number of NLRIs in the update sent: max 4, min 0
(output omitted)

```

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

```
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 the specified RD.**

© 2006 Cisco Systems, Inc. All rights reserved. MPLS v2.2—5-14

The **show ip bgp vpnv4** command displays IP version 4 (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

```
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 ?
```

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2--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`

```
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
```

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—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 a VRF 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-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

This 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 Cisco Express Forwarding (CEF) and label forwarding information base (LFIB) structures.

Monitoring per-VRF CEF and LFIB Structures

```
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

© 2006 Cisco Systems, Inc. All rights reserved. MPLS v2.2—5-17

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

- The **show ip cef vrf** command displays the VRF Forwarding Information Base (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.)

```
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 Se1/0.2, 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 Se1/0.2, point2point, tags imposed: {26 39}
```

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

© 2006 Cisco Systems, Inc. All rights reserved.MPLS v2.2--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*]].

The label stack in the VRF table can be inspected using the **show ip cef vrf vrf-name detail** command. The *tags imposed* values in the output displays the MPLS label stack. The first label in the MPLS label stack is the Label Distribution Protocol (LDP) label forwarded toward the egress provider edge (PE) router, and the second label is the VPN label advertised by the egress PE router.

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, Packet over SONET (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.)

```
Router#show mpls forwarding vrf SiteA2
```

Local tag	Outgoing tag or VC	Prefix or Tunnel Id	Bytes tag switched	Outgoing interface	Next Hop
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 tag	Outgoing tag or VC	Prefix or Tunnel Id	Bytes tag switched	Outgoing interface	Next Hop
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

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2--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.

- **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) Displays IP prefix address (in dotted decimal notation) and length of mask (0 to 32)
<i>mask</i>	(Optional) Displays 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

Router#

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

- Displays labels associated with VPNv4 routes

```
Router#show ip bgp vpnv4 all labels

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

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—5-20

You can use the **show ip bgp vpnv4 all labels** command to display tags assigned to local or remote VRF routes by the local or remote PE router. This command displays labels 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

```
Router#  
telnet host /vrf vrf-name
```

- Performs PE-CE Telnet through specified VRF

```
Router#  
ping vrf vrf-name ip-address
```

- Performs ping based on VRF routing table

```
Router#  
trace vrf vrf-name ip-address
```

- Performs VRF-based traceroute

© 2006 Cisco Systems, Inc. All rights reserved. MPLS v2.2--5-21

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

- The **telnet** command can be used to connect to a customer edge (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 that were discussed in this lesson.

Summary

- **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 *vrf-name*
 - show ip route vrf *vrf-name*
 - show ip bgp vpnv4 vrf *vrf-name*
- **Use the show ip bgp neighbors command to monitor MP-BGP sessions.**

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—5-22

Summary (Cont.)

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

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—5-23

Configuring OSPF as the Routing Protocol Between PE and CE Routers

Overview

This lesson explains the provider edge-customer edge (PE-CE) routing protocol configuration steps required when you are running Open Shortest Path First (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 provider edge (PE) and customer edge (CE) routers when you are setting up Multiprotocol Label Switching (MPLS) Virtual Private Networks (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 and 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
- 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
- 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 these 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 Border Gateway Protocol (BGP) backbone into the customer network. Isolated copies of the Interior Gateway Protocol (IGP) run at every site, and Multiprotocol Border Gateway Protocol (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



© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—5-5

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 implementation, any route that is redistributed into OSPF from another routing protocol will now be redistributed as an *external* OSPF route. The OSPF routes received by one PE router are propagated across the MPLS backbone and redistributed back into OSPF at another site as external OSPF routes.

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

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

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

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—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.


This list 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 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

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


© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—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



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 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.

Note 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

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.

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—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.

Note	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

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

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—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 multi-exit discriminator (MED) attribute.

OSPF Superbackbone: Implementation (Cont.)



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 or type 3 LSA is always generated because the receiving PE router acts as an ABR between the superbackbone and the OSPF area (or areas).

Note If the OSPF superbackbone connects two or more instances of the same area, the redistributed route will appear as an interarea summary LSA.

OSPF Superbackbone: External Routes



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


© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—5-13

The external OSPF routes are redistributed into MP-BGP in exactly the same way as the internal OSPF routes. Here is 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



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

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

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.**

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—5-15

To configure OSPF as a PE-CE routing protocol, you need to start a separate OSPF process for each virtual routing and forwarding instance (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.)

```
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.**

© 2006 Cisco Systems, Inc. All rights reserved.MPLS v2.2—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.

Note

Prior to Cisco IOS Release 12.3(4)T, 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 IGP. The number of PE-CE routing processes was therefore limited to 28.

This restriction was removed for MPLS in Cisco IOS Release 12.3(4)T.

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.)

```
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.**


© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—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


- **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.**

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2--5-19

These 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 discussed in the “Introducing the MPLS VPN Routing Model” lesson of the “MPLS VPN 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.



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 Multiprotocol Internal Border Gateway Protocol (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.)

- 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.

© 2006 Cisco Systems, Inc. All rights reserved.


MPLS v2.2—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.)



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: Routing Loops Across OSPF Domains

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.
2	The down bit is set because the route should not be redistributed back into MP-BGP.
3	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.
4	The OSPF route is propagated through the other OSPF domain with the down bit cleared.
5	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

- **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.**

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—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, these 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 autonomous system (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

- 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.


© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—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



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.
2	The tag field is set to the BGP AS number, and the down bit is set. 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	The route is propagated across the OSPF domain with the tag set to the BGP AS number.
5	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.)

- **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.**

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—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.)

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.

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—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.)




© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—5-31

Because the sham link is seen as an intra-area link between PE routers, an OSPF adjacency is created and a database exchange (for the particular OSPF process) occurs across the link. The PE router can then flood LS100

As between sites from across the MPLS VPN backbone. As a result, the desired intra-area connectivity is created.



Because the OSPF cost of the sham link has been configured as preferred over the backdoor link, the PE routers will prefer routes learned via the high-bandwidth backbone. The implementation results in optimum packet flow.

Configuring a Sham Link

This topic describes how to configure a sham link.

Configuring a Sham Link

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

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—5-33

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

The criteria listed here 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.)

```
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.

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2--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</i> [<i>mask</i>]
<i>destination-address</i>	IP address of the destination PE router in the format: <i>ip-address</i> [<i>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.



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 that were discussed in this lesson.

Summary

- **OSPF areas connect to a common backbone area in a two-tier hierarchical model.**
- **Basic OSPF across an MPLS VPN includes a BGP backbone. OSPF is run at each site, while MP-BGP is used to propagate routes between each site.**
- **A better option implements the MP-BGP backbone as a new transparent OSPF superbackbone above existing areas.**
- **OSPF PE-CE routing is implemented as a separate routing process. (One routing process per VRF.)**

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—5-36

Summary (Cont.)

- **A sham link is the OSPF down bit prevents routing loops.**
- **A sham link is the OSPF tag field is also used to prevent routing loops.**
- **Packet forwarding is optimized across the MPLS VPN using the OSPF routing bit**
- **A sham link is required between any two VPN sites that belong to the same OSPF area and share an OSPF backdoor link.**
- **The area sham-link cost command is used to configure a sham link across a MPLS VPN backbone.**

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—5-37

Configuring BGP as the Routing Protocol Between PE and CE Routers

Overview

This lesson explains the provider edge-customer edge (PE-CE) routing protocol configuration steps required when you are using Border Gateway Protocol (BGP) as the routing protocol between provider edge (PE) and customer edge (CE) routers.

It is important to understand not only what you can configure between PE and CE routers when you are setting up Multiprotocol Label Switching (MPLS) Virtual Private Networks (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:

- 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
- Implement SOO for loop prevention

Configuring a Per-VRF BGP Routing Context

This topic describes how to configure a per-virtual routing and forwarding (VRF) BGP routing context.

Configuring per-VRF BGP Routing Context

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.**

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—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 Interior Gateway Protocol (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 IP version 4 (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.)



The PE router can be defined as a BGP neighbor.

Example: Configuring per-VRF BGP Routing Context

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

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

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

Limiting the Number of Routes in a VRF

- **SPs offering MPLS VPN services are at risk of denial-of-service attacks similar to those aimed at SPs 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.

© 2006 Cisco Systems, Inc. All rights reserved. MPLS v2.2—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 (SP) 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 these 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 Multiprotocol Internal Border Gateway Protocol (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

```
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)

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—5-6

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

- **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).

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—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.)

```
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* option creates a syslog error message when the maximum number of routes exceeds the threshold.
- Syslog messages generated by this command are rate-limited.

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—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>	This parameter 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>	This parameter rejects routes when the threshold limit is reached. The threshold limit is a percentage of the limit specified, from 1 to 100.
warn-only	This parameter 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.



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 autonomous system (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

- **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.**

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—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.)

- **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.**

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2--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 farther down the AS path of the AS number of the receiving router are not replaced because they indicate a real routing information loop.
- An extra copy of the sending router AS number is prepended to the AS path. The standard AS number prepending procedure occurs on every External Border Gateway Protocol (EBGP) update.

AS-Override: Command

```
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.

© 2006 Cisco Systems, Inc. All rights reserved.MPLS v2.2—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 VPN version 4 (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 (65213) with the provider AS number (65115) 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.)



- **VPN perspective:** VPN-A is connected to VPN-B via CE-AB.
- **Physical topology:** The CE-AB router is dual-connected to the PE routers.
- **MPLS VPN perspective:** The CE-AB router has two links into the P-network.
- **BGP perspective shows issue:** The CE-AB router has two connections to AS 65115.

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—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 provider network (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 issue appears when the setup is analyzed through the BGP perspective:

- The CE (CE-AB) router has to propagate routes between two PE routers, which are both in the same AS.

Allowas-in: The Issue (Cont.)



- PE-1 announces network 10.1.0.0/16 to CE-AB.
- CE-AB 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-AB, which may require a Cisco IOS software upgrade on the CE router.

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2--5-18

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).

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

Example: Using Allowas-in to Support Customer Site Linking Two VPNs

This example discusses using the **allowas-in** command on the PE router to support a customer site linking two VPNs. This example is similar to the situation in which two customer sites use the same AS number.

Allowas-in: Implementation

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.

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—5-19

Networks may 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 so that incoming updates that exceed the limit can be dropped.

Allowas-in: Command

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

```
neighbor allowas-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.

© 2006 Cisco Systems, Inc. All rights reserved.MPLS v2.2--5-20

neighbor allowas-in

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

- **neighbor allowas-in** *number*
- **no neighbor allowas-in** *number*

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

Syntax Description


Parameter	Description
<i>number</i>	This parameter 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 Site of Origin (SOO) for loop prevention.



Most aspects of BGP loop prevention are bypassed when either the AS-override feature or the allows-in feature is used. 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 allows-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.)

- **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.**

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2--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.)

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

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—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 route target (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 these combinations: ■ <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 Multiprotocol Border Gateway Protocol (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.)

Other Inbound Routing Updates

Router(config-if)#

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

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

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—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.

Note	An example configuring SOO for EIGRP PE-CE loop prevention was discussed in the “Configuring EIGRP PE-CE Routing” topic of the “Configuring Small-Scale Routing Protocols Between PE and CE Routers” lesson.

Implementing SOO for Loop Prevention (Cont.)

```
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

© 2006 Cisco Systems, Inc. All rights reserved.MPLS v2.2—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 that were discussed in this lesson.

Summary

- **Use the address-family ipv4 vrf *vrf-name* command in the BGP routing process to configure a per-VRF BGP routing context.**
- **SPs 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 routes command to limit the total number of VRF routes.**

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—5-26

Summary (Cont.)

- **BGP loop detection prevents customers from reusing their AS number. The neighbor *ip-address* as-override command prevents this issue by replacing the customer AS number with the ISP AS number.**
- **By default, a customer site cannot link two VPN sites of the same AS number because of BGP loop detection. The neighbor allowas-in *number* command 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.**

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—5-27

Troubleshooting MPLS VPNs

Overview

This lesson explains the preliminary steps for troubleshooting a Multiprotocol Label Switching (MPLS) Virtual Private Network (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 LSP
- 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

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)?

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—5-3

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

- Is Cisco Express Forwarding (CEF) enabled on all routers in the transit path between the provider edge (PE) routers?
- Are labels for Border Gateway Protocol (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

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 IPv4 routes propagated to other CE routers?

© 2006 Cisco Systems, Inc. All rights reserved.


MPLS v2.2—5-4

Verification of the routing information flow should be done systematically, starting at the ingress customer edge (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



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.**

© 2006 Cisco Systems, Inc. All rights reserved. MPLS v2.2—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. 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

The diagram illustrates a network topology for validating PE-to-PE routing information flow. It features a central box labeled 'P-Network' containing two Provider Edge (PE) routers, PE-1 and PE-2. PE-1 is connected to two CE-Spoke routers on its left, and PE-2 is connected to two CE-Spoke routers on its right. A curved arrow indicates a loopback path on PE-1. A straight arrow points from the left CE-Spoke routers towards PE-1. Another straight arrow points from PE-2 towards the right CE-Spoke routers. The diagram is enclosed in a rectangular frame.

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.**


© 2006 Cisco Systems, Inc. All rights reserved. MPLS v2.2—5-6

The CE routes received by the PE router need to be redistributed into Multiprotocol BGP (MP-BGP); otherwise, the routes 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-virtual routing and forwarding (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 route distinguisher (RD) prepended to the IPv4 prefix and the route targets (RTs) attached to the CE route can be verified with the **show ip bgp vpnv4 vrf vrf-name ip-prefix** command.

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



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.**

© 2006 Cisco Systems, Inc. All rights reserved.


MPLS v2.2—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.)



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.**


© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—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 multi-exit discriminator (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.)



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.**

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2-5-9


The VPN version 4 (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.)



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.

© 2006 Cisco Systems, Inc. All rights reserved.


MPLS v2.2—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) course for more information on route redistribution troubleshooting.

Validating PE-to-CE Routing Information Flow

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



Last but not least, the routes redistributed into the PE-CE routing protocol have to be propagated to CE routers. You may also configure the CE routers with a default route toward the PE routers. 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

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?**

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—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

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.**

© 2006 Cisco Systems, Inc. All rights reserved. MPLS v2.2—5-13

One of the most common configuration mistakes related to data flow is the failure to enable CEF on 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 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

```
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
```

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2--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



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.

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2--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 mpls forwarding-table** 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 mpls forwarding-table** 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

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.

© 2006 Cisco Systems, Inc. All rights reserved. MPLS v2.2—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 Interior Gateway Protocol (IGP) and maximum transmission unit (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** configuration command on the P and PE routers. The **traceroute** command from the ingress PE router toward the BGP next hop should display no intermediate hops when TTL propagation is disabled. If intermediate hops are displayed, the 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 label forwarding information base (LFIB) status.

Validating the LFIB Status

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.

© 2006 Cisco Systems, Inc. All rights reserved. MPLS v2.2--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 that were discussed in this lesson.

Summary

- **Divide MPLS troubleshooting into two main steps:**
 - **Verify routing information flow.**
 - **Verify proper data flow.**
- **Validate CE-to-PE routing information flow by checking 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.**
- **Verify that routes are redistributed back into the CE routing protocol on the PE route and propagated toward CE routers to validate PE-to-CE routing information flow.**

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—5-18

Summary (Cont.)

- **Verify data flow systematically, starting at the ingress 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, review the contents of the LFIB on the egress PE router in comparison to the second label in the label stack on the ingress PE router.**

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—5-19

Module Summary

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

Module Summary

- **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, define the 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.**

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—5-1

Module Summary (Cont.)

- **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.**

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—5-2

An MPLS VPN implementation involves 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.

References

For additional information, refer to these resources:

- Access Cisco.com for additional information about VPNs.
- Cisco Systems, Inc. The “BGP Filtering and Route Selection” module in the *Configuring BGP on Cisco Routers* (BGP) course.
- Cisco Systems, Inc. The “Advanced BGP Configuration” module in the BGP course.
- Cisco Systems, Inc. *Building Scalable Cisco Internetworks* (BSCI) 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) VRFs are assigned to an interface. (Source: Using MPLS VPN Mechanisms of Cisco IOS Platforms)
- A) true
 - B) false
- 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 existing IP address of an interface 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 labels 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: _____

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? (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 SOO? (Source: Configuring BGP as the Routing Protocol Between PE and CE Routers)

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

Q43) Where can you set the SOO? (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) 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 (VPNv4) is used to carry VPN route updates between PE routers.

Note	IBGP (IPv4) is used to carry <i>non-VPN</i> 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) Enable the static route using the **ip route vrf name static route parameters** command.
- Q20) Use 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-table vrf
- Q27) **show ip bgp vpnv4 all labels**
- 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**
- 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 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. 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, 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 vpnv4 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) Determine 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

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 Multiprotocol Label Switching (MPLS) Virtual Private Network (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
- Configure selective VRF imports
- Configure selective VRF exports

What Are Advanced VRF Features?

This topic identifies the advanced VRF features.

Advanced VRF Features

- Selective import**
 - This features allows you to specify additional criteria for importing routes into the VRF.
- Selective export**
 - This features allows you to specify additional RTs attached to exported routes.
- VRF route limit**
 - This features allows you to specify the maximum number of routes in a VRF to prevent memory exhaustion on PE routers or denial-of-service attacks.

© 2006 Cisco Systems, Inc. All rights reserved. MPLS v2.2—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

- **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.**

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—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.)

```
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.

© 2006 Cisco Systems, Inc. All rights reserved.MPLS v2.2—5-5

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 version 4 (VPNv4) routes and to import into the VRF only routes matching a certain prefix.



A configuration similar to this one could be used to accomplish the following:

- Deploy advanced MPLS VPN topologies (for example, a managed router services topology)
- Increase the security of an extranet VPN by allowing only predefined subnetworks to be inserted into a VRF, thus preventing an extranet site from inserting unapproved subnetworks into the extranet

Note A similar function is usually not needed in an intranet scenario because all customer routers in an intranet are usually under common administration.

Configuring Selective VRF Export

This topic describes how to configure selective VRF export.

Configuring Selective VRF Export

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 a VRF:

- This route map can set extended community RTs.
- No other set operations can be performed by this route map.

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—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.)

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.

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—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: ■ <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 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.)

```
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.

© 2006 Cisco Systems, Inc. All rights reserved. MPLS v2.2—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.



Note Depending on when you configure the **export map** command, you may need to use the **clear ip bgp** command to force the existing BGP session to propagate the extended communities.

In the earlier example, selective import of routes into a VRF was achieved with an import route map in the receiving VRF that allowed only routes from a certain address block to be inserted into the VRF. In this example, routes from a certain address block are marked with an additional RT in the originating VRF and are automatically inserted into the receiving VRF on the basis of their RT.

The main difference between import and export route maps is therefore the deployment point:

- The import route map is deployed in the receiving VRF.
- The export route map is deployed in the originating VRF.

Based on the network design, the functionality of one or the other 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 that were discussed in this lesson.

Summary

- **Two advanced VRF features were discussed:**
 - **Selective import**
 - **Selective export**
- **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.**

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—6-11

Introducing Overlapping VPNs

Overview

Overlapping Virtual Private Networks (VPNs) are usually used to connect parts of two separate VPNs. A third VPN is created within the Multiprotocol Label Switching (MPLS) VPN network that contains sites from both VPNs. A new route target (RT) extended community is used for networks originating in the sites that are also in the new VPN. This action may require a new virtual routing and forwarding (VRF) instance, 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
- 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 interconnection, two simple VPNs are created, each containing a customer central site and its remote sites. Then a third VPN, which partially overlaps with the customer VPNs but connects only their central sites, is created. The central sites can talk to each other. Each central site can also talk to the remote sites in its simple VPN, but not to the remote sites belonging to the other customer simple VPN. The addresses used in the central sites, however, 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

- **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**

© 2006 Cisco Systems, Inc. All rights reserved. MPLS v2.2—6-4


These are two typical uses for overlapping VPNs:

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

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

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:


- Export all networks with RTs 123:750 *and* 123:1001.
- Import 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:

- Export all networks with RTs 123:760 *and* 123:1001.
- Import 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:

- The simple VPN for customer A contains routes that originate from the following:
 - A-Central site
 - A remote sites
- The simple VPN for customer B contains routes that originate from the following:
 - B-Central site
 - B remote sites
- The overlapping VPN contains routes that originate from the following:
 - A-Central site
 - B-Central site
- All 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 and RTs.)
- 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 and RTs.)
- A-Central is a member of VPN-A and overlapping VPN-AB. (It cannot use the same RD as A-Spoke-1 and A-Spoke-2 because it has different routing requirements. It needs a unique RD and RTs for VPN-A as well as Central.)
- B-Central is a member of VPN-B and overlapping VPN-AB. (It needs a unique RD as well as RTs for VPN-B as well as 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



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 PE routers.

Summary

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

Summary

- **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 the required RTs based on the VPN membership.**

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—6-9

Introducing Central Services VPNs

Overview

A central services Virtual Private Network (VPN) is used when multiple 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. The other VPNs, however, are not able to see one another. The central services VPN is implemented using two route target (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 other 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
- 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 these characteristics:


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

This topology can be used in these situations:

- The service provider offers services to all customers by allowing them access to a common VPN.
- Two (or more) companies want to exchange information by sharing a common set of servers.
- A security-conscious company separates its departments and allows them access only to common servers.

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:

- Client 1 and client 2 have their own RTs (123:101, 123:102) that they import and export; they also export networks with RT 123:193 and import networks with RT 123:192.

Note Client-specific RTs were introduced to comply with the implementation requirements of Cisco IOS Software Release 12.0 T, in which each virtual routing and forwarding (VRF) instance 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 123:193 and exports networks with RT 123:192.

Client 1 does the following:

- Exports all networks with RTs 123:101 or 123:193
- Imports all networks that carry RT 123:101 or 123:192

Client 2 does the following:


- Exports all networks with RTs 123:102 or 123:193
- Imports all networks that carry RT 123:102 or 123:192

The central site does the following:

- Exports all networks with RT 123:192 or RT 123:199
- Imports all networks that carry RT 123:193 or RT 123:199

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 customer edge (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 these requirements:

- You need a separate VRF for each client.
- You need one VRF per provider edge (PE) router connecting a server site.
- You need a unique route discriminator (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.)



Server sites:

- Use one VRF for each service type.
- Use a unique RD on each 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.

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—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
		123:192	123:193
Client-2	123:102	123:102	123:102
		123:192	123:193

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	123:104	123:104	123:104
		123:192	123:193
Client-5	123:105	123:105	123:105
		123:192	123:193

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
		123:193	123:192

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
		123:193	123:192




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.)



- 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.


© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—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 Central Services and Simple VPNs

This topic identifies the RD requirements when you are integrating a central services VPN with a simple VPN.




For this design, you need two RDs per VPN:

- Configure one RD for simple VPN sites. (The same value should also be used for import and export RTs.)
- Configure one RD for the central services VRFs.

Identifying the RT Requirements When Integrating Central Services and Simple VPNs

This topic identifies the RT requirements when you are integrating a central services VPN with a simple VPN.



This table shows an RD and RT numbering scheme for PE-1.

PE-1 RD and RT Numbering Scheme

VRF	RD	Import RT	Export RT
VPN-A	123:750	123:750	123:750
VPN-B	123:760	123:760	123:760
VPN-A-Central	123:751	123:750 123:101	123:750 123:100

This table shows an RD and RT numbering scheme for PE-2.

PE-2 RD and RT Numbering Scheme


VRF	RD	Import RT	Export RT
VPN-A	123:750	123:750	123:750
VPN-B	123:760	123:760	123:760
VPN-B-Central	123:761	123:760	123:760
		123:101	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	123:101
		123:100	

Configuring VRFs in a Central Services VPN and Simple VPN



```
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
```

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—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 using the RD and RT numbering scheme presented in the tables.

Summary

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

Summary

- **A central services VPN is used to provide access from centralized servers to one or more customers.**
- **A central services VPN routing model indicates these requirements:**
 - 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 these requirements:**
 - 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 these:**
 - Use a separate VRF for each client.
 - Use a unique RD on each client site.
 - Use a unique RD in each set of server sites.
 - Use import and export RT matching between server and client sites.

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—6-14

Summary (Cont.)

- **The hybrid of a simple VPN and a central VPN provides the following:**
 - Customers have 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. 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.**

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—6-15

Introducing the Managed CE Routers Service

Overview

A service provider can use a separate network management Virtual Private Network (VPN) to manage the customer edge (CE) routers of all the VPNs. A pair of route target (RT) extended communities is used to accomplish this goal. One RT exports CE router loopback addresses and is imported into the virtual routing and forwarding (VRF) instance 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 VRF instances. 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
- 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 route discriminator (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

- **Configure the per-customer import-export RT in all customer VRFs.**
- **Configure the NMS import-export RT in NMS VRF.**
- **Import routes with the NMS RT into the customer VRF.**
- **Export loopback addresses from the customer VRF with RT NMS_Client.**
- **Import routes with RT NMS_Client into NMS VRF.**

© 2006 Cisco Systems, Inc. All rights reserved. MPLS v2.2-6-5

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	123:750
		123:101	123:100 (using NMS route-map)
VPN-B	123:760	123:760	123:760
		123:101	123:100 (using NMS route-map)

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



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 provider edge (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 that were discussed in this lesson.

Summary

- 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 a central and simple VPN, except for the RT marking process during route export.

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—6.7

Module Summary

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

Module Summary

- **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.**

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—6-1

Market forces continually drive service providers to provide more complex centralized services for their customers. These services, such as advanced virtual routing and forwarding (VRF) import and export features, overlapping Virtual Private Networks (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.
- Cisco Systems, Inc. “NAT Integration with MPLS VPNs.”
- Cisco Systems, Inc. “DHCP Relay—MPLS VPN Support.”
- Cisco Systems, Inc. *Multicast VPN Design Guide*

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 RDs? How many RTs? (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 RDs do you need for a central services VPN solution with 50 client sites and 3 server sites? How many RTs? (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)

Module Self-Check Answer Key

- Q1) A selective VRF import command allows you to select routes to be imported into a VRF based on criteria other than the VRF RT.
- Q2) The import route map filter is combined with the RT import filter—a route has to pass the RT import filter first and then pass the import route map to be imported into the VRF.
- Q3) A selective VRF export command allows you to attach specific RTs to a subset of routes exported from a VRF. (By default, the same RTs get attached to all exported routes.)
- Q4) A route map can be specified for each VRF to attach additional RTs to routes exported from a VRF. The export route map performs only the attachment of RTs; it does not perform any filtering function, and you cannot change any other route attributes with it.
- Q5) extended community RTs
- Q6) companies that use MPLS VPNs to implement both intranet and extranet services, or a security-conscious company that wishes to limit visibility between different departments in the organization
- Q7) Selected sites in a VPN can communicate only with sites within their VPN. Other selected sites can communicate with sites in their VPN and selected sites in a second VPN.
- Q8) VRFs: 3
route distinguishers: 4
route targets: 3
- Q9) in solutions where some sites (server sites) can communicate with all other sites, but all the other sites (client sites) can communicate only with the server sites
- Q10) The clients have their own RTs that they import and export; the clients also export networks with an RT that will be used by the services VRF and import networks with an RT that identifies the routes of the services site. The services site imports and exports networks with the RT of its VPN, but it also imports networks with RTs that identify the client sites.
- Q11) Client PE router:
ip vrf Client_1
rd 123:101
route-target both 123:101
route-target export 123:303
route-target import 123:203
ip vrf Client_2
rd 123:102
route-target both 123:102
route-target export 123:303
route-target import 123:203
- Server PE router:
ip vrf Server
rd 123:103
route-target both 123: 203
route-target import 123:303
- Q12) route distinguishers = 51 route targets = 52
You need one RD for each client (50) and one RD (1) shared by both server sites. You need one RT for each client (50) to export its routes to its VPN, one RT (1) for all of the clients to export their routes to the server, and one RT (1) that is shared by both server sites to export their routes to the clients.
- Q13) Create a simple VPN that provides connectivity for all of the customer sites that do not need access to the central services.
Next, create a simple VPN that provides access between all of the server sites that are in the service.
Finally, create an overlapping VPN that contains the sites that must have access to both the customer VPNs and the services VPN.

- Q14) If the service provider is managing the customer routers, it is convenient to have a central point that has access to all CE routers but not to the other destinations at customer sites.
- 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

Internet Access and MPLS VPNs

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 Internet connectivity scenarios
- Describe implementing separate Internet access and MPLS VPN services
- Describe implementing Internet access as a separate VPN service on a MPLS VPN backbone

Introducing Internet Access with MPLS VPNs

Overview

This lesson describes common customer Internet connectivity scenarios and identifies two design models for combining Internet access with Multiprotocol Label Switching (MPLS) Virtual Private Network (VPN) services. The lesson lists the benefits and drawbacks of these models, and then gives an overview of 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:


- Describe common customer Internet connectivity scenarios
- Identify the two major design models for combining Internet access with MPLS VPN services
- Describe the benefits and drawback of Internet access through global routing
- Describe the benefits and drawback of Internet access in a separate VPN
- Describe the disadvantages of providing Internet access through route leaking

Customer Internet Connectivity Scenarios

This topic describes some customer Internet connectivity scenarios.

Classical Internet Access

Classical Internet access is implemented through a central firewall that connects the customer network to the Internet in a secure fashion.




The customer network and the Internet are connected only through the firewall. The addressing requirements for this type of connection are 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 customer's private addresses and the public addresses assigned to the customer by the Internet service provider. Alternatively, the firewall might perform an application-level proxy function that also isolates private and public IP addresses.

There are a number of benefits associated with this design. It is a well-known setup, and the 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 that needs to be managed.

The major drawback of this design is the traffic flow. All traffic from the customer network to the Internet passes through the central firewall. Although this might not be a drawback for smaller customers, it can be a severe limitation for large organizations with many users, especially when they are geographically separated.

Multisite Internet Access

Some customers find the traffic flow limitations of the central firewall setup too limiting. To bypass the limitations of Internet access through a central firewall, some customers are using designs in which each customer site has its own independent Internet access.




This design solves traffic flow issues, but the associated drawback is higher exposure. Each site has to be individually secured against unauthorized Internet access, which leads to the increased complexity of managing a firewall at every customer site.

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 but different logical links to minimize costs.

For customers that do not want the complexity of managing their own firewalls, a managed firewall service offered by the service provider can help address the security issues of Internet connectivity.

Wholesale Internet Access


A service provider may provide wholesale Internet access to customers from a range of upstream Internet service providers (ISPs) to satisfy the connectivity and reliability requirements of various customers.



The selection of upstream ISPs and the corresponding configuration processes should therefore be as easy as possible for the service provider. From an Internet perspective, customers A, B, and C are connected to ISP X or ISP Y. This means that the IP address space used by the customer should be allocated from the block of addresses administered by the selected ISP. The service provider providing wholesale Internet access may have to use a different address for every upstream ISP.

Internet Design Models for Service Providers

This topic identifies two major design models for service providers combining Internet access with MPLS VPN services.



Because Internet access is one of the most popular services that service providers offer their customers, many service providers offer Internet access as well as MPLS VPN service on their shared backbone. Integrating Internet access with an MPLS VPN solution is one of the most common service provider business requirements. A background of common customer Internet connectivity scenarios will help in assessing possible implementations.

Major Design Models

Network designers who want to offer Internet access and MPLS VPN services in the same backbone can choose between these two major design models:

- Internet access that is implemented through global routing on the provider edge (PE) routers and that is not a VPN service
- Internet access that is implemented as a separate VPN in the ISP network

Major Design Models

Two major design models:

- **Internet access separate from VPN services**
- **Internet access as a separate VPN**

© 2006 Cisco Systems, Inc. All rights reserved. MPLS v2.2—7-7

In both cases, security should be the most important concern for customers when they connect to public networks. Customers should isolate private VPNs from Internet traffic, either physically on a separate interface or on a subinterface. Appropriate firewall support either in a dedicated device or integrated in the router Cisco IOS software is a necessity. Depending on the network addressing, network address translation (NAT) will be needed for most customers.

Internet Access Through Global Routing

This topic describes options for implementing Internet access through global routing.

Internet Access Through Global Routing

- **Implementation via separate interfaces that are not placed in any VRF, via either:**
 - Static default routing on a PE
 - BGP between CE and PE
- **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

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—7-8

Implementing Internet access through global routing is identical to building a traditional IP backbone offering Internet services. Depending on whether or not the customer sites need full Internet routes, static routes or Border Gateway Protocol (BGP) is used for routing to the Internet. For full Internet routes, BGP is deployed between the PE routers and the PE Internet gateway to exchange Internet routes, and the global routing table on the PE routers is used to forward the customer traffic toward Internet destinations. The PE routers may or may not have the full Internet routing table. The PEs and the provider Internet gateway are in the same Internal BGP (IBGP) area.

The PE routers also use Multiprotocol BGP (MP-BGP) to support VPNs for their customers. The customers reach the global routing table by using a separate logical link for Internet access. Internet access through global routing on separate logical links is easy to set up, 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 encapsulation that supports subinterfaces (for example, Frame Relay).

Internet Access as a Separate VPN

This topic describes the benefits and drawbacks of having Internet access in a separate VPN.

Internet Access Through a Separate VPN Service

- **Implementation through a separate VPN**
- **Benefit:**
 - The provider backbone is isolated from the Internet; increased security is realized.
- **Drawback:**
 - All Internet routes are carried as VPN routes; full Internet routing cannot be implemented because of scalability problems.

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—7.9

For a service provider, implementing Internet access through a separate VPN is similar to offering another managed VPN service.

The major benefit of implementing Internet access as a separate VPN is increased isolation between the provider backbone and the Internet—which results in increased security for the provider. The flexibility of MPLS VPN topologies also provides for some innovative design options that allow service providers to offer services that were simply 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 having all the Internet routes inside a single VPN.

Disadvantages of Providing Internet Access Through Route Leaking

A third variant for Internet support with MPLS VPNs is to provision route leaking into the global routing context. Although this not a recommended design, this option is briefly discussed to show an alternate practice that has been used in the industry.

Internet Access Through Route Leaking

- **Implementation through corporate VPN**
- **Benefit:**
 - Does not use a separate connection for Internet traffic
- **Drawback:**
 - Insecure because Internet traffic is mingled with corporate traffic in the VPN
 - Harder to apply security policies on mingled traffic
 - Cannot implement full Internet routing because of scalability problems

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2--7-10

Some customers may want to obtain Internet access across their corporate VPN by leaking routes between the VRF and global routing tables.

Note	For security reasons, this approach is not recommended. It is not a good practice to bring in Internet traffic using the corporate VPN. This practice negates the isolation of the corporate VPN.

With route leaking, the customer site uses a static default route in the virtual routing and forwarding (VRF) table pointing to the global next-hop address of an Internet gateway. Any packets that use the default route leave the VPN space and are routed based on the global routing table at the PE that is the next-hop router. This feature allows “leaking” of VPN packets into the global address space.

The next-hop address for the static default route points to the interface where the Internet routes are learned on the Internet gateway. This interface address must be present within the global routing table so that label switching of the packets between the PE and Internet gateway router can occur. This means that the Internet gateway router must advertise its Internet interface address within the Interior Gateway Protocol (IGP) that runs across the service provider backbone.

Because there is no redistribution between IP version 4 (IPv4) and VPN version 4 (VPNv4) routes, a mechanism is needed to allow the PE router to resolve the next-hop address for the Internet gateway from within the VRF. This is achieved through the use of the **global** keyword within the static default route configuration. The **global** keyword specifies that the next-hop address of the static route is resolved within the global routing table, not within the VRF.

Each of the customer IP subnets that are used as source addresses for Internet access must be advertised through the global BGP version 4 (BGP-4) process toward the Internet. The easiest solution for placing these routes in the global routing table is to configure further global static routes so that the VPN subnets appear within the global routing table as well as the VRF. This makes them available for advertisement via BGP-4. These static routes must point to an interface, with a next-hop address specified if the outbound interface is multiaccess (such as Ethernet). Then either these routes can be redistributed into BGP or the **network** command can be used within the BGP process. If redistribution is used, a route map can be utilized to specify which addresses to advertise.

Note	This approach is not recommended. It is not discussed further in this course. This option is briefly discussed only to show an alternate practice that has been used in the industry.

Summary

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

Summary

- **Classical Internet access connects through a central firewall. You can use a centralized ISP managed firewall service.**
- **Multisite Internet access connects the firewall of every site. You can use a centralized ISP-managed firewall service.**
- **Wholesale Internet access service offers connectivity to multiple ISPs.**

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—7-11

Summary (Cont.)

- **There are two recommended service provider designs for combining Internet access with MPLS VPN services:**
 - **Global routing (Internet access not from a VPN), which uses separate interfaces that are not placed in any VRF**
 - **Internet services as a separate VPN, which allows for service provider separation of backbone and Internet traffic**
- **Route leaking is insecure and not recommended because of this approach negates isolation of the corporate VPN.**

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—7-12

Implementing Separate Internet Access and VPN Services

Overview

This lesson describes implementing Internet access service totally separate from Multiprotocol Label Switching (MPLS) Virtual Private Network (VPN) services. It is important to understand why you may choose to use global routing to separate Internet access from VPN services.

This lesson identifies the provider edge-customer edge (PE-CE) requirements for separating Internet access from VPN services and describes 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 the features of classical Internet access for a VPN customer
- Describe how separate subinterfaces are implemented to support Internet access using global routing
- Describe how Internet access can be obtained for every customer site
- Identify the benefits and limitations of separating Internet access from VPN services

Classical Internet Access for a VPN Customer

This topic describes the features of classical Internet access for a VPN customer on a shared service provider backbone.



The classical Internet access design for a customer is based on a separate Internet access model. One central customer site has connectivity to the Internet and provides access to the rest of the customer sites. The central site either connects through a firewall or runs the Cisco IOS firewall feature set.

In the shared service provider backbone, the provider edge (PE) routers have full or partial Internet routes to offer the customer. The provider's Internet gateway (Internet-GW) is in the same Internal Border Gateway Protocol (IBGP) domain as the provider (P) and PE routers.

This design model can easily map to a customer with an MPLS VPN implementation. In this example, the customer network has been interconnected with an MPLS VPN. A central customer edge router (CE-Central) has Internet connectivity and provides Internet access through a firewall for all the sites in the customer network.

This traditional Internet access implementation model provides maximum design flexibility, because the Internet access is completely separated from the MPLS VPN services. However, the limitations of traditional IP routing prevent this implementation method from being used for innovative Internet access solutions, such as wholesale Internet access.

Using Separate Subinterfaces

For security reasons, the VPN and Internet routes to a customer should be on separate interfaces. This topic describes how separate subinterfaces can be used for implementing Internet access through global routing.

Using Separate Subinterfaces

- **Separate physical links for VPN and Internet traffic are sometimes not acceptable because of high cost.**
- **Subinterfaces could be used.**
 - Over WAN links using Frame Relay or ATM encapsulation (including xDSL)
 - Over LAN links
- **A tunnel interface could be used.**
 - Over a VRF-aware tunnel, so that VPN traffic does not run over a global tunnel

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—7-4


Instead of separate physical links for VPN and Internet traffic, subinterfaces can be used to create two logical links over a single physical link. Subinterfaces can be configured only on WAN links 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 other encapsulation types, a tunnel interface can be used between the CE router and the PE router. Depending on the router platform and Cisco IOS version, virtual routing and forwarding (VRF)-aware tunnels are now supported.

- VRF-aware tunnels remove the need for the endpoints of the tunnel to be in a global routing table.
- Without a VRF-aware tunnel, MPLS VPN traffic would need to be tunneled across the Internet interface.

Note Further information on VRF-aware tunnels is outside the scope of this course.

Example Configuration: Static Routes



Example: Internet Access Through Static Routes

Using static routing on the CE and the PE routers is the simplest and most common implementation for providing Internet access. The figure illustrates a configuration used to implement Internet access through two Frame Relay subinterfaces with static routes. In this simple example, CE-Central does not need to receive full Internet routing; it just needs a default static route to PE3 to reach the Internet. PE3 has a route to the Internet through Internet-GW, as well as a static route for the Customer A subnets pointing to CE-Central. The full Internet routing table needs to be present only on Internet-GW.


The following configuration steps are performed:

- The customer VRF instance VPN-A is created.
- The VPN subinterface (Serial0.1) is created and associated with data-link connection identifier (DLCI) 101. The subinterface is added to the VPN-A VRF.
- The CE router is configured as a BGP neighbor inside the VPN in the VRF VPN-A with the **neighbor ip-address remote-as as-number** command in address-family ipv4 configuration mode.
- A separate Internet subinterface (Serial0.2) is created and associated with DLCI 102.
- A static route to the public CE routes is created and placed in the global routing table.

In addition to the PE configuration, the CE router implements a default static route pointing to 171.68.10.1 on PE3.

Note On the CE-Central router, distribution of the default route may be needed so that remote sites can also access the Internet. Issues of security and private addresses would have to be resolved.

Example Configuration: Dynamic Routes




Example: Dynamic Internet Access Through a Separate Subinterface

The figure illustrates the configuration needed to implement Internet access through a dedicated Frame Relay interface. Some additional configuration steps are needed, compared to the previous static example:

- The CE router is configured as a BGP neighbor in the global BGP process with the **neighbor ip-address remote-as as-number** command in router configuration mode.
- A static route to the public CE routes is not needed, because the central CE will advertise the company routes to the PE through the BGP process

In this example, PE3 configures CE-Central as a BGP peer in the global BGP process using the non-VPN interface as the neighbor address. CE-Central will also need to establish peering with PE3 and announce the customer routes that it needs to advertise.

Note On the CE routers, static routes or redistribution may need to be implemented so that appropriate customer sites can access the Internet. Issues of security and private addresses would have to be resolved.




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.

Accessing the Internet from Every Customer Site

This topic describes how Internet access is implemented at every customer site.



Another option is to provide separate Internet access at every customer site. In this case, two physical (or logical) links between every CE router and its PE router would be needed. This design often becomes too complex or too expensive to implement. Issues including customer route propagation to the Internet and securing access at multiple access points would need to be resolved.

Note	The allowas-in feature may need to be configured on the PE router if the customer is propagating individual site routes to the Internet through BGP.

Separate Internet Access Benefits and Limitations

This topic describes the benefits and limitations of separate Internet access.

Benefits and Limitations of Separate Internet Access for the Service Provider

Benefits:

- **Well-known model**
- **Supports all customer requirements**
- **Allows all Internet services implementations, 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.**

© 2006 Cisco Systems, Inc. All rights reserved. MPLS v2.2—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.
- This model allows all Internet services implementations, 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 that were discussed in this lesson.

Summary

- **Classical Internet access for a VPN customer is based on a separated Internet access design model**
- **Separate subinterfaces can be used for implementing Internet access through global routing**
- **Internet access from every customer site can be supported but is often too complex or too expensive with classic Internet access.**
- **The main drawback of separate Internet access is that PE routers potentially carry full Internet routing table**

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—7-19

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 Virtual Private Network (VPN). The lesson identifies the scaling issues of this design, and discusses how to implement the design in a Multiprotocol Label Switching (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 the features of a redundant Internet access implementation
- Configure classical Internet access for a VPN customer
- Configure Internet access from every customer site
- Describe how to implement wholesale Internet access
- Identify the benefits and limitations of running an Internet backbone in a VPN

Internet Access as a Separate VPN

This topic describes the features of using a separate MPLS VPN to provide Internet access.

Internet Access as a Separate VPN

- **A provider Internet gateway is connected as a CE router to the MPLS VPN backbone.**
- **The Internet gateway does not insert full Internet routing into the Internet VPN.**
 - **Only the default route and the local (regional) routes are inserted.**
- **Every customer site that needs Internet access is assigned to the same Internet VPN as the Internet gateway.**


© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—7.3

The MPLS VPN architecture can provision a separate VPN to provide Internet access for VPN customers. The service provider defines the Internet VPN and can use different MPLS VPN topologies to implement various types of Internet access. Under this design model, the provider Internet gateways appear as customer edge (CE) routers to the MPLS VPN backbone. Customer Internet access is enabled using a dual VPN topology supporting both an Internet VPN and a customer VPN across separate customer interfaces.

In this design, the Internet VPN should not contain the full set of global Internet routes, because that would make the solution completely nonscalable. The provider Internet gateway routers should announce a default route toward the provider edge (PE) routers. To optimize local routing, the local and regional Internet routes should be inserted in the Internet VPN.

Internet Access as a Separate VPN (Cont.)



- The Internet VPN is isolated from the P routers.


© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—7-4

When the service provider implements Internet access as a separate VPN, the Internet backbone is carried on a VPN, which is isolated from the provider backbone. This topology results in increased security for the provider backbone, because Internet hosts can reach only PE routers, not the core provider routers (P routers). The VPN customers are connected to the Internet simply through an additional virtual routing and forwarding (VRF) instance at the PE.

Example: Configuring the Internet Gateway in a Separate VPN

This section shows an example configuration implementing the Internet gateway in a separate VPN.




The figure shows a sample configuration of the Internet-GW router with conditional default route advertisement. Router Internet-GW will advertise the default route to the PE-GW router only if Internet-GW can reach the network 192.168.0.0/16.

The following steps are used to configure this functionality:

- Step 1** A static default route is configured toward a next hop in network 192.168.0.0. If the network 192.168.0.0 is not reachable, this static route will not enter the IP routing table.
- Step 2** The default route origination is configured in the Border Gateway Protocol (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 192.168.0.0/16 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.
- Step 4** PE-GW configures the Internet VRF and applies it to the appropriate interface.

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:

- 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 Internet-GW router as the exit point toward those destinations.
- The MED attribute can also be used to indicate the preferred default route to the PE routers. In this setup, one Internet-GW router acts as a primary Internet gateway, and the other Internet-GW router acts as a backup.
- The redundancy established so far covers the path between customer sites and the Internet-GW routers. A failure in the Internet backbone might break the Internet connectivity for the customers if the Internet-GW routers announce the default route unconditionally. Conditional advertisement of the default route is therefore configured on the Internet-GW routers—the Internet-GW routers announce the default route to the PE routers only if the Internet-GW routers can reach an upstream destination.

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 VPN over the MPLS VPN backbone. The link between a PE router and the Internet-GW router is assigned to the Internet VRF as discussed before. Internet-GW announces a default route to the Internet.

One link between the PE router and each central customer router is assigned to the customer VRF, and one is assigned to the Internet VRF.

In this example, PE3 connects CE-Central of customer A to both the VPN-A VRF and the Internet VRF. A prefix list is used on PE3 in the Internet address family configuration to filter BGP advertisements to the CE-Central router.

Classical Internet Access for a VPN Customer (Cont.)




The CE-Central router has two neighbor relationships with PE3. It announces the summary route for range of networks that have been assigned to customer A. These networks may be at remote locations.

In this example, the remote locations CE-A1 and CE-A2 will need a default route across VPN-A to reach CE-Central and through CE-Central reach the Internet.

If needed, an External BGP (EBGP) multihop session can be configured between the Internet gateway (Internet-GW) and the CE-Central router to give full Internet routes to the customer. These routes would not be injected into the Internet VRF.

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 can be implemented by configuring the Internet VRF at every location. This solution adds complexity for the customer, because firewall and NAT support may be needed at every site unless a central managed firewall service is offered by the service provider.

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 Internet service provider (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

Benefits:

- Supports all Internet access service types
- Can support all customer requirements, including a BGP session with the customer, accomplished through advanced BGP setup

Drawbacks:

- Full Internet routing cannot be carried in the VPN; default routes are needed that can lead to suboptimal routing.
- Internet gateway routers act as CE routers on the VPN backbone; implementing overlapping Internet and VPN backbones requires care.

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—7-11

Internet access implemented as a separate VPN has the following benefits:

- This design model supports all Internet access services, ranging from traditional Internet access to innovative services such as wholesale Internet access.
- This design also supports all customer requirements, including full Internet routing on customer routers through an EBGp multihop session with the Internet gateway (Internet-GW).

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.
- The Internet backbone gateway router is positioned as a customer edge route connected to the MPLS VPN backbone. If the service provider runs Internet service and MPLS VPN service on the same set of routers, the interconnection between the two services requires special considerations.

The benefits of the separate VPN design far outweigh the limitations.

Note An MPLS VPN extension is carrier supporting carrier (CSC). The CSC feature enables one MPLS VPN-based service provider to allow other service providers to use a segment of its backbone network. Discussion of the CSC feature is not within the scope of this course.

Summary

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

Summary

- **MPLS VPN architecture supports defining the Internet as a VPN.**
 - Redundant Internet access is easy to achieve.
 - The classical Internet access model can be easily implemented using the Internet VPN.
- **Internet access from every customer site can be implemented by configuring the Internet VRF on a second interface at every location**
- **Wholesale Internet access can be implemented by creating a separate VPN for every upstream ISP.**
- **Internet VPNs supports all customer requirements, including full Internet routing.**

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2--7-12

Module Summary

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

Module Summary

- **Separating Internet access from VPN services can be accomplished through use of logical subinterfaces.**
- **Classical Internet access uses a separate logical IPv4 interface.**
- **MPLS VPN architecture can implement the Internet as another VPN.**

© 2006 Cisco Systems, Inc. All rights reserved. MPLS v2.2—7-1

There are several different models that are used to combine Internet access with Multiprotocol Label Switching (MPLS) Virtual Private Network (VPN) services. Each model has benefits and drawbacks. It is important to understand the implications of implementing the various models.

References

For additional information, refer to this resource:

- Access Cisco.com for additional information about MPLS Internet access.
- MPLS Internet Connectivity Options
http://www.cisco.com/en/US/tech/tk436/tk428/technologies_white_paper09186a00801281f1.shtml

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 Internet Access Models with MPLS VPNs)
- A) All of the customer Internet traffic passes through the central firewall service.
 - B) None of the customer Internet traffic passes through the central firewall service.
 - C) Only some of the customer Internet traffic passes through the central firewall service.
 - D) There is no drawback.
- Q2) When the wholesale Internet access solution is implemented, which statement is correct? (Source: Introducing Internet Access Models with MPLS VPNs)
- 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.
- Q3) What are the two major design models for implementing Internet access and MPLS VPNs? (Choose two.) (Source: Introducing Internet Access Models with MPLS VPNs)
- A) using a separate VPN for Internet access
 - B) using global route leaking
 - C) using global routing on PE routers
 - D) using default routes
- Q4) What are two major benefits of using VPNs to provide Internet access? (Choose two.) (Source: Introducing Internet Access Models with MPLS VPNs)
- A) Only some of the customer Internet traffic needs to pass through a firewall.
 - B) The provider backbone is isolated from the Internet.
 - C) Security is increased.
 - D) The Internet VRF can support full Internet routes.
- Q5) What is the recommended implementation option for using global routing to provide Internet access? (Source: Introducing Internet Access Models with MPLS VPNs)
- A) The separate subinterfaces are placed in separate VRFs.
 - B) The Internet subinterface is not placed in a VRF.
 - C) One interface for the Internet and the MPLS VPN is recommended.
 - D) Static routes are always needed for Internet access.

- Q6) Which is the main benefit of using separate interfaces or subinterfaces to provide Internet access? (Source: Introducing Internet Access Models with MPLS VPNs)
- A) Wholesale Internet access or central firewall service cannot be implemented.
 - B) It is easy to implement internal BGP sessions between the PE router and CE router.
 - C) They are well-known and are easy to implement.
 - D) The separate subinterfaces are placed in separate VRFs.
- Q7) The traditional Internet access implementation model provides _____. (Source: Implementing Separate Internet Access and VPN Services)
- A) minimum design flexibility, because the Internet access is completely separated from the MPLS VPN services
 - B) maximum design flexibility, because the Internet access is completely separated from the MPLS VPN services
 - C) the ability to put the separate subinterfaces into separate VRFs
 - D) no ability to use wholesale Internet services
- Q8) In situations where cost prohibits separate physical links for VPN and Internet traffic for classical Internet access, _____. (Source: Implementing Separate Internet Access and VPN Services)
- A) the Internet access should be integrated with the MPLS VPN through route leaking
 - B) the Internet access should be integrated with the MPLS VPN through route peaking
 - C) VRFs can be used to create two logical links over a single physical link
 - D) subinterfaces can be used to create two logical links over a single physical link
- Q9) Which setup for a VPN customer is based on a separated Internet access design model? (Source: Implementing Separate Internet Access and VPN Services)
- A) classical Internet access
 - B) route peaking
 - C) route leaking
 - D) classical Intranet access
- Q10) For customers that need Internet access from every site, two physical (or logical) links between every CE router and its PE router might be _____. (Source: Implementing Separate Internet Access and VPN Services)
- A) too complex to implement
 - B) too expensive to implement
 - C) too complex and too expensive to implement
 - D) the most secure solution to implement
- Q11) The Internet VPN should _____. (Source: Implementing Internet Access as a Separate VPN)
- A) contain the full set of Internet routes because that would supportable optimal routing
 - B) not contain the full set of Internet routes because that would make the solution completely nonscalable
 - C) support route leaking because that would make the solution completely scalable
 - D) not be used for optimal scalability

- Q12) What should multiple provider Internet gateways in the Internet VPN advertise to the PE routers? (Source: Implementing Internet Access as a Separate VPN)
- A) the full set of Internet routes, because that would support optimal routing
 - B) any leaked routes, to support optimal routing
 - C) the full set of Internet routes, because that would provide redundant routing
 - D) the default route, for routing redundancy
- Q13) In a classical Internet access using the separate VPN model, the link between a PE router and the Internet gateway 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)
- A) Internet, customer
 - B) customer, Internet
 - C) Internet, Internet
 - D) customer, customer
- Q14) The main disadvantage of having Internet access at every customer site is that _____. (Source: Implementing Internet Access as a Separate VPN)
- A) this option cannot be implemented with a VPN solution
 - B) firewall services would be needed at every site
 - C) any leaked routes support suboptimal routing
 - D) wholesale Internet access cannot be implemented
- Q15) A central managed firewall service can be implemented by _____. (Source: Implementing Internet Access as a Separate VPN)
- A) creating a separate Internet VPN for every customer
 - B) assigning customers to their ISPs by default routes in the Internet VRF
 - C) creating a separate customer VPN
 - D) implementing the managed firewall service through the Internet gateway
- Q16) Wholesale Internet access can be implemented by _____. (Source: Implementing Internet Access as a Separate VPN)
- A) creating a separate VPN for every upstream ISP
 - B) assigning customers to their ISPs by default routes in the Internet VRF
 - C) creating a separate VPN for every customer
 - D) propagating full Internet routes into the Internet VRF
- Q17) Which two are drawbacks of Internet access that is implemented as a separate VPN? (Choose two.) (Source: Implementing Internet Access as a Separate VPN)
- A) Full Internet routing cannot be carried inside a VRF.
 - B) The customer has no ability to receive full Internet routes.
 - C) Default routing toward the Internet gateways has to be used, which can potentially result in suboptimal routing.
 - D) The customer must use default routes.

Module Self-Check Answer Key

- Q1) A
- Q2) B
- Q3) A, C
- Q4) B, C
- Q5) B
- Q6) C
- Q7) B
- Q8) D
- Q9) A
- Q10) C
- Q11) B
- Q12) D
- Q13) C
- Q14) B
- Q15) D
- Q16) A
- Q17) A, C

MPLS TE Overview

Overview

This module provides a brief overview of Multiprotocol Label Switching (MPLS) traffic engineering (TE). The module covers the requirement for TE in modern networks and presents the basic concepts and mechanics that support TE, including tunnel path discovery with link-state protocols and tunnel path signaling with Resource Reservation Protocol (RSVP). The module describes the basic MPLS TE commands for the implementation of MPLS traffic tunnels.

Module Objectives

Upon completing this module, you will be able to describe the tasks and commands that are necessary to implement MPLS TE. This ability includes being able to meet these objectives:

- Identify basic TE and MPLS TE concepts
- Identify MPLS TE components
- Create and configure MPLS TE
- Monitor basic MPLS TE on Cisco IOS platforms

Introducing the TE Concept

Overview

This lesson describes the concepts that allow service providers and enterprises to map traffic through specific routes to optimize network resources, especially bandwidth. Traffic engineering (TE) enables backbone networks to be engineered to deliver the total subscribed capacity to customers more efficiently.

This lesson is a helpful introduction for learners who are planning to improve the usage of their network resources with Multiprotocol Label Switching (MPLS) TE.

Objectives

Upon completing this lesson, you will be able to describe the basic TE concepts. This ability includes being able to meet these objectives:

- Identify the concepts behind TE
- Identify the major business drivers for implementing TE
- Identify congestion avoidance and describe how TE can reduce some congestion-avoidance issues
- Identify how TE is implemented using a Layer 2 overlay model
- Identify how TE is implemented using a Layer 3 model
- Identify how TE is implemented using the MPLS TE model

What Is TE?

This topic identifies the underlying concepts of TE.

What Is Traffic Engineering?

- **Traffic engineering is manipulating your traffic to fit your network.**
- **Network engineering is building your network to carry your predicted traffic.**
- **TE is commonly used in voice telephony networks.**
- **TE is a process of measures, models, and controls of traffic to achieve various goals.**
- **TE for data networks provides an integrated approach to managing traffic at Layer 3.**

© 2006 Cisco Systems, Inc. All rights reserved. MPLS v2.2—3.3

Traffic engineering (TE) can be contrasted to network engineering:

- Traffic engineering is manipulating your traffic to fit your network
- Network engineering is building your network to carry your predicted traffic

TE has been widely used in voice telephony. TE means that the traffic is measured and analyzed. A statistical model is then applied to the traffic pattern to make a prognosis and estimates.

If the anticipated traffic pattern does not match the network resources well, the network administrator remodels the traffic pattern. Such decisions can be made to achieve a more optimal use of the resources or to reduce costs by selecting a cheaper transit carrier.

In the data communications world, traffic engineering provides an integrated approach to engineering traffic at Layer 3 in the Open Systems Interconnection (OSI) model. The integrated approach means that routers are configured to divert traffic from destination-based forwarding and move the traffic load from congested parts of the network to uncongested parts.

Traditionally, this diversion has been done using overlay networks where routers use carefully engineered ATM permanent virtual circuits (PVCs) or Frame Relay PVCs to distribute the traffic load on Layer 2.

What Is Traffic Engineering? Traffic Engineering Motivations

- **Reduce the overall cost of operations by more efficient use of bandwidth resources**
- **Prevent a situation where some parts of a network are overutilized (congested), while other parts remain underutilized**
- **Implement traffic protection against failures**
- **Enhance SLA in combination with QoS**

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—8-4

Cost reduction is the main motivation for TE.

A cost savings, which results from a more efficient use of resources, will help to reduce the overall cost of operations.

Additionally, more efficient use of bandwidth resources means that a service provider or enterprise network manager can avoid a situation where some parts of a network are congested, while other parts are underutilized.

Because TE can be used to control traffic flows, it can also be used to provide protection against link or node failures by providing backup tunnels.

Finally, when combined with quality of service (QoS) functionality, TE can provide enhanced service level agreements (SLAs).

Business Drivers for TE

This topic identifies some business drivers that may lead to implementation of TE.

Business Drivers for Traffic Engineering

- **Routers forward traffic along the least-cost route discovered by routing protocols.**
- **Network bandwidth may not be efficiently utilized:**
 - The least-cost route may not be the only possible route.
 - The least-cost route may not have enough resources to carry all the traffic.
 - Alternate paths may be underutilized.

© 2006 Cisco Systems, Inc. All rights reserved. MPLS v2.2—8-5

In a Layer 3 routing network, packets are forwarded hop by hop. In each hop, the destination address of the packet is used to make a routing table lookup. The routing tables are created by an Interior Gateway Protocol (IGP), which finds the least-cost route, according to its metric, to each destination in the network.

In many networks, this method works well. But in some networks, the destination-based forwarding results in the overutilization of some links, while others are underutilized. This imbalance will be the case when there are several possible routes to reach a certain destination. The IGP selects one of them as the best and uses only that route. In the extreme case, the best path may have to carry so large a volume of traffic that packets are dropped, while the next-best path is almost idle.

One solution to the problem would be to adjust the link bandwidths to more appropriate values. The network administrator could reduce the bandwidth on the underutilized link and increase the bandwidth on the overutilized one. However, making this adjustment is not always possible.

The alternate path is a backup path. In the case of a primary link failure, the backup must be able to forward at least the major part of the traffic volume that is normally forwarded by the primary path. Therefore, it may not be possible to reduce the bandwidth on the backup path. And without a cost savings, the budget may not allow an increase to the primary link bandwidth.

To provide better network performance within the budget, network administrators move a portion of the traffic volume from the overutilized link to the underutilized link. During normal operations, this move results in fewer packet drops and quicker throughput. In the case of a failure to any of the links, all traffic is forwarded over the remaining link, which then, of course, becomes overutilized.

Moving portions of the traffic volume cannot be achieved by traditional hop-by-hop routing using an IGP for path determination.

Business Drivers for Traffic Engineering (Cont.)

- **Lack of resources results in congestion in two ways:**
 - **When network resources themselves are insufficient to accommodate offered load**
 - **When traffic streams are inefficiently mapped onto available resources**
- **Some resources are overutilized while others remain underutilized.**

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—8-6

Network congestion, caused by too much traffic and too few network resources, cannot be solved by moving portions of the traffic between different links. Moving the traffic will help only in the case where some resources are overutilized and others are underutilized. The traffic streams in normal Layer 3 routing are inefficiently mapped onto the available resources.

Good mapping of the traffic streams onto the resources constitutes better use of the money invested in the network.

Cost savings that result in a more efficient use of bandwidth resources help to reduce the overall cost of operations. These reductions, in turn, help service providers and organizations gain an advantage over their competitors. This advantage becomes more and more important as the service provider market gets more and more competitive.

A more efficient use of bandwidth resources means that a provider could avoid a situation where some parts of its network are congested while other parts are underutilized.

Congestion Avoidance and TE

This topic discusses congestion avoidance and how TE can reduce some congestion issues.

Congestion Avoidance and Traffic Engineering

Network congestion can be addressed by either:

- **Expansion of capacity or classical congestion control techniques (queuing, rate limiting, and so on)**
- **Traffic engineering, if the problems result from inefficient resource allocation**

The focus of TE is not on congestion created as a result of a short-term burst, but on congestion problems that are prolonged.

© 2006 Cisco Systems, Inc. All rights reserved.


MPLS v2.2—5-7

TE does not solve temporary network congestion that is caused by traffic bursts. This type of problem is better handled by an expansion of capacity or by classical techniques such as various queuing algorithms, rate limiting, and intelligent packet dropping. TE does not solve problems when the network resources themselves are insufficient to accommodate the required load.

TE is used when the problems result from inefficient mapping of traffic streams onto the network resources. In such networks, one part of the network suffers from prolonged congestion, possibly continuously, while other parts of the network have spare capacity.

TE with a Layer 2 Overlay Model

This topic discusses the use of a Layer 2 overlay model with TE.



In the Layer 2 overlay model, the routers (Layer 3 devices) are overlaid on top of the Layer 2 topology. The routers are not aware of the physical structure and the bandwidth that is available on the links. The IGP views the PVCs or switched virtual circuits (SVCs) as point-to-point links and makes its forwarding decisions accordingly.


All engineering is done at Layer 2. PVCs are carefully engineered across the network, normally using an offline management system. SVCs are automatically established using signaling, and their way across the Layer 2 network is controlled by integrated path determination, such as the Private Network-to-Network Interface (PNNI) protocol.

In the Layer 2 overlay model, PVCs or SVCs carry the traffic across the network. With a Frame Relay network, PVC setup is most often made using a management tool that helps the network administrator calculate the optimum path across the Layer 2 network with respect to available bandwidth and other constraints that may be applied on individual links.

ATM may use the same type of tools as Frame Relay for PVC establishment or may use the SVC approach, where routers use a signaling protocol to dynamically establish an SVC.

If the Layer 2 network provides a full mesh between all routers, the Layer 3 IGP sees all the other routers as directly connected and, most likely, uses the direct logical link whenever it forwards a packet to another router. The full mesh gives Layer 2 full control of the traffic load distribution. Manual engineering of PVCs and the configuration of PNNI parameters are the tools that allow very exact control of how the traffic uses the available bandwidth.

Traffic Engineering with a Layer 2 Overlay Model: Example



Traffic engineering in Layer 2, using the overlay model, allows detailed decisions about which link should be used to carry various traffic patterns.

In this example, traffic from R2 to R3 uses the red PVC (solid arrows), which takes the shortest path using the upper transit switch.

However, traffic from R1 to R3 uses the blue PVC (dashed arrows), which does not take the shortest path. TE on Layer 2 has been applied to let the second PVC use links that would otherwise have been underutilized. This approach avoids overutilization of the upper path.

Traffic Engineering with a Layer 2 Overlay Model (Cont.)

Drawbacks of the Layer 2 overlay solution

- **Extra network devices**
- **More complex network management:**
 - Two-level network without integrated network management
 - Additional training, technical support, field engineering
- **IGP routing scalability issue for meshes**
- **Additional bandwidth overhead (“cell tax”)**
- **No differential service (class of service)**

© 2006 Cisco Systems, Inc. All rights reserved.


MPLS v2.2--8-10

Using the Layer 2 overlay model has several drawbacks:

- The routers are not physically connected to other routers. The Layer 2 network introduces the need for an additional device, the ATM or Frame Relay switch.
- Two networks must be managed. The Layer 2 network requires its own management tools, which, among several other tasks, support TE as well. At the same time, the router network (Layer 3), with its IGP and tuning parameters, must be managed. Both these management tasks require trained staff for technical support and in the field.
- The Layer 3 network must be highly meshed to take advantage of the benefits that are provided by the Layer 2 network. The highly meshed network may cause scalability problems for the IGP because of the large number of neighbors.
- Overlay networks always require an extra layer of encapsulation. A Frame Relay header must be added to the IP packets, or, when ATM is used, the IP packet must be segmented into cells, each of which must have its own header. The extra layer of encapsulation causes bandwidth overhead.
- The Layer 2 devices do not have any Layer 3 knowledge. After the router has transmitted the IP packet across the physical link to the first switch, all IP knowledge is lost. When congestion does occur in the Layer 2 network, the switches have no ability to selectively discard IP packets or to requeue them. Thus, no IP differentiated services can be used within the Layer 2 switch network.

TE with a Layer 3 Model

This topic discusses the use of a Layer 3 model with TE.



If the same network topology is created using routers (Layer 3 devices), TE must be performed differently. In the example here, if no traffic engineering is applied to the network, traffic from both R8 and R1 toward R5 will use the least-cost path (the upper path). This flow may result in the overutilization of the path R2, R3, R4, R5, while the path R2, R6, R7, R4, R5 will be underutilized.

Traffic Engineering with a Layer 3 Model (Cont.)

The current forwarding paradigm, centered around “destination-based,” is clearly inadequate:

- **Path computation based just on IGP metric is not enough.**
- **Support for “explicit” routing (source routing) is not available.**
- **Supported workarounds are static routes, policy routing.**
- **Provide controlled backup and recovery.**

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2--8-12


The destination-based forwarding paradigm that is currently used in Layer 3 networks cannot resolve the problem of overutilization of one path while an alternate path is underutilized.

The IGP uses its metric to compute a single best way to reach each destination. There are problems with Layer 3 TE:

- IP source routing could be used to override the IGP-created routing table in each of the intermediate routers. However, in a service provider network, source routing is most often prohibited. The source routing would also require the host to create the IP packets to request source routing. The conclusion is that source routing is not an available tool for TE.
- Static routing, which overrides the IGP, can be used to direct traffic to take a different path than that of other traffic. However, static routing does not discriminate among various traffic flows based on the source. Static routing also restricts how redundancy in the network can be used.
- Policy-based routing (PBR) is able to discriminate packet flows based on the source, but it suffers from low scalability and the same static routing restrictions as to how redundancy can be used.

TE with the MPLS TE Model

This topic discusses using TE with the MPLS TE model.



In the MPLS TE implementation, routers use MPLS label switching with TE.

The aim is to control the paths along which data flows, rather than relying simply on destination-based routing. MPLS TE uses tunnels to control the data flow path. An MPLS TE tunnel is simply a collection of data flows that share some common attribute. This attribute might be all traffic sharing the same entry point to the network and the same exit point.

A TE tunnel maps onto an MPLS label-switched path (LSP). After the data flows and the TE tunnels are defined, MPLS technology is used to forward traffic across the network. Data is assigned an MPLS TE, which defines the route taken through the network. The packets that are forwarded under MPLS TE have a stack of two labels imposed by the ingress router. The topmost label identifies a specific LSP or TE tunnel to use to reach another router at the other end of the tunnel. The second label indicates what the router at the far end of the tunnel should do with the packet.

Note Additional details on MPLS tunnels are covered in the next lesson.

Traffic Engineering with the MPLS TE Model (Cont.)

- **The MPLS TE LSPs are created by RSVP.**
- **The actual path can be specified:**
 - Explicitly defined by the system administrator
 - Dynamically defined using the underlying IGP protocol

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2--8-14

For MPLS TE, manual assignment and configuration of the labels can be used to create LSPs to tunnel the packets across the network on the desired path. However, to increase scalability, the Resource Reservation Protocol (RSVP) is used to automate the procedure.

By selecting the appropriate LSP, a network administrator can direct traffic via explicitly indicated routers. The explicit path across identified routers provides benefits that are similar to those of the overlay model without introducing a Layer 2 network. This approach also eliminates the risk of running into IGP scalability problems because of the many neighbors that exist in a full mesh of routers.

MPLS TE provides mechanisms equivalent to those described previously in this lesson in connection with the Layer 2 overlay network. For circuit-style forwarding, instead of using ATM or Frame Relay virtual circuits, the MPLS TE tunnel is used. For signaling, RSVP is used with various extensions to set up the MPLS TE tunnels.

For constraint-based routing (CBR) used in MPLS TE, either Intermediate System-to-Intermediate System (IS-IS) or Open Shortest Path First (OSPF) with extensions is used to carry resource information, such as available bandwidth on the link. Both link-state protocols use new attributes to describe the nature of each link with respect to the constraints. A link that does not have the required resource is not included in the MPLS TE tunnel.

To actually direct the traffic onto the MPLS TE tunnels, network administrators need extensions to IS-IS and OSPF. Directing the traffic into tunnels results in the addition of entries in the Forwarding Information Base (FIB). The IP packets are directed into the MPLS TE tunnel by imposing the correct label stack.

Summary

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

Summary

Traffic engineering measures, models, and controls traffic to achieve various goals.

- TE is driven by inefficient bandwidth utilization.
- TE focuses on prolonged congestion problems.
- With the TE Layer 2 overlay model, routers are not aware of the physical structure and bandwidth available on links.
- With the TE Layer 3 model, the destination-based forwarding paradigm cannot handle the problem of overutilization of one path while an alternate path is underutilized.
- TE with the MPLS TE model means that the routers use the MPLS label-switching paradigm.

© 2006 Cisco Systems, Inc. All rights reserved. MPLS v2.2—8-16

References

For additional information, refer to this resource:

- RFC 3209, *RSVP-TE: Extensions to RSVP for LSP Tunnels*

Understanding MPLS TE Components

Overview

This lesson explains the components of Multiprotocol Label Switching (MPLS) traffic engineering (TE), such as traffic tunnels (along with associated characteristics and attributes), tunnel path discovery based on link-state protocols, and tunnel setup signaling with Resource Reservation Protocol (RSVP).

Objectives

Upon completing this lesson, you will be able to describe the basic components of MPLS TE. This ability includes being able to meet these objectives:

- Identify, at a conceptual level, how a traffic tunnel functions
- Identify traffic tunnel characteristics
- Identify traffic tunnel attributes
- Identify the relation between network links and link attributes
- Identify the function of constraint-based path computation
- Identify TE procedures
- Identify the role of RSVP in path setup and admission control
- Identify how using TE modifies the forwarding table mechanisms

Traffic Tunnels: Concepts

This topic describes the concept of traffic tunnels.

Traffic Tunnels: Concepts

The concept of MPLS TE traffic tunnels was introduced to overcome the limitations of hop-by-hop IP routing:

- **A tunnel is an aggregation of traffic flows that are placed inside a common MPLS label-switched path.**
- **Flows are then forwarded along a common path within a network.**

© 2006 Cisco Systems, Inc. All rights reserved.


MPLS v2.2—5-3

The aim of TE is to control the paths along which data flows, rather than relying simply on “normal” destination-based routing. To fulfill this aim, the concept of a “traffic tunnel” has been introduced.

A traffic tunnel is simply a collection of data flows that share some common attribute:

- Most simply, this attribute might be the sharing of the same entry point to the network and the same exit point. In practice, this point could be an Internet service provider (ISP) network, where there is a definable data flow from the points of presence (POPs), where the customers attach to the ISP network. There are also the Internet exchange points (IXPs), where data typically leaves the ISP network to traverse the Internet.
- In a more complex situation, this attribute could be augmented by defining separate tunnels for different classes of service. For example, in an ISP model, leased-line corporate customers could be given a preferential throughput over dial-in home users. This preference might be greater guaranteed bandwidth or lower latency and higher precedence. Even though the traffic enters and leaves the ISP network at the same points, different characteristics could be assigned to these types of users by defining separate traffic tunnels for their data.

Traffic Tunnels: Concepts (Cont.)



- **Unidirectional single class of service model encapsulates all of the traffic between an ingress and an egress router.**
- **Different classes of service model assigns traffic into separate tunnels with different characteristics.**

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—8-4

Defining traffic trunks (tunnels) requires an understanding of the traffic flows in the network. By understanding the ingress and corresponding egress points, a picture of the traffic flows in the network can be produced.

In the example, there are two traffic tunnels (TT1 and TT2) that are defined for data from PE1 to PE3. These tunnels are unidirectional; they identify the traffic flows *from* PE1.

Note In practice, there are probably similar tunnels operating in the opposite direction, *to* PE1.

There may also be tunnels that are defined from all the other routers to each other. Defining tunnels from every router in the network to every router might sound like an administrative nightmare. However, this is not usually the case for the following reasons:

- The routers that are identified are on the edge of the network. The traffic tunnels link these routers across the core of the network.
- In most networks it is relatively easy to identify the traffic flows, and they rarely form a complete “any-to-any” mesh.
- For example, in ISP networks, the traffic tunnels generally form a number of star formations, with their centers at the IXPs and the points at the POPs. Traffic in an ISP network generally flows from the customers that are connected at the POPs to the rest of the Internet (reached via the IXPs). A star-like formation can also exist in many networks centering on the data center. This tendency is found in both ISP networks (providing web-hosting services) and enterprise networks.

After the data flows, and therefore the traffic tunnels, are defined, the technology that they use to forward the data across the network is MPLS. Data that enters a traffic tunnel is assigned an MPLS label-switched path (LSP). The LSP defines the route that is taken through the network.

Traffic Tunnels: Characteristics

This topic describes the characteristics of traffic tunnels.

Traffic Tunnels – Characteristics

- **A traffic tunnel is distinct from the MPLS LSP through which it traverses:**
 - **More than one TE tunnel can be defined between two points:**
 - Each tunnel may pick the same or different paths through the network
 - Each tunnel will use different MPLS labels
 - **A traffic tunnel can be moved from one path onto another based on resources in the network.**
- **A traffic tunnel is configured by defining its required attributes and characteristics.**

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—5-5


Traffic tunnels are distinct from the MPLS LSPs that they use in two key ways:

- There is a one-to-one mapping of traffic tunnels onto MPLS LSPs. Two tunnels may be defined between two points and may happen to pick the same path through the network. However, they will use different MPLS labels.
- Traffic tunnels are not necessarily bound to a particular path through the network. As resources change in the core, or perhaps as links fail, the traffic tunnel may reroute, picking up a new MPLS LSP as it does.

Configuring the traffic tunnels includes defining the characteristics and attributes that it requires. In fact, defining the characteristics and attributes of traffic tunnels is probably the most important aspect of TE. Without a specification of the requirements of the data in a traffic tunnel, the data might as well be left to route “normally” based only on destination information over the least-cost path.

Traffic Tunnels: Attributes

This topic describes the attributes of traffic tunnels.



A traffic tunnel is a set of data flows sharing some common feature, attribute, or requirement. If there is no characteristic in the data flow in common with some other flow, there is nothing to define that data as part of a flow or group of flows.

Therefore, the traffic tunnel must include attributes that define the commonality between the data flows making up the tunnel. The attributes that characterize a traffic tunnel include the following:

- **Ingress and egress points:** These points are, fundamentally, the routers at the ends of the tunnel. They are the most basic level of commonality of data flows, given that the flows in a tunnel all start in the same place and end in the same place.
- **Complex characteristics of the data flows:** Examples are bandwidth and latency and precedence requirements.
- **Class of data:** This attribute defines what data is part of this tunnel and what is not. This definition includes such characteristics as traffic flow, class of service, and application class.

The network administrator defines the attributes of a traffic tunnel when the tunnel itself is defined. However, some of these attributes are, in part, influenced by the underlying network and protocols.

Note MPLS TE setup is a control plane function.

Traffic Tunnels—Attributes (Cont.)

The administrator enters the relevant information (attributes) at the headend of the traffic tunnel:

- **Traffic parameter**—Resources required for tunnel (for example, required bandwidth)
- **Generic path selection and management**—Path can be administratively specified or computed by the IGP
- **Resource class affinity**—Include or exclude certain links for certain traffic tunnels
- **Adaptability**—Should the traffic tunnel be reoptimized?
- **Priority and preemption**—Importance of a traffic tunnel and possibility for a preemption of another tunnel
- **Resilience**—Desired behavior under fault conditions

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—8-7

The general tunnel characteristics must be configured by the network administrator to create the tunnel. This configuration includes some or all of the following:

- **Traffic parameters:** Traffic parameters are the resources that are required by the tunnel, such as the minimum required bandwidth.
- **Generic path selection and management:** This category refers to the path selection criteria. The actual path that is chosen through the network could be statically configured by the administrator or could be assigned dynamically by the network, based on information from the Interior Gateway Protocol (IGP), which is Intermediate System-Intermediate System (IS-IS) or Open Shortest Path First (OSPF).
- **Resource class affinity:** This category refers to restricting the choice of paths by allowing the dynamic path to choose only certain links in the network.

Note This restriction can also be accomplished by using the IP address exclusion feature.

- **Adaptability:** Adaptability is the ability of the path to reroute on failure or to optimize on recovery or discovery of a better path.
- **Priority and preemption:** Traffic tunnels can be assigned a priority (0 to 7) that signifies their “importance.” When you are setting up a new tunnel or rerouting, a higher-priority tunnel can tear down (preempt) a lower-priority tunnel; in addition, a new tunnel of lower priority may fail to set up because some tunnels of a higher priority already occupy the required bandwidth of the lower-priority tunnel.
- **Resilience:** Resilience refers to how a traffic tunnel responds to a failure in the network. Does it attempt to reroute around failures or not?

Network Links and Link Attributes

This topic discusses the relationship between network links and link attributes.

Network Links and Link Attributes

Resource attributes (link availability) are configured locally on the router interfaces:

- **Maximum bandwidth**
 - The amount of bandwidth available
- **Link affinity string**
 - To allow the operator to administratively include or exclude links in path calculations
- **Constraint-based specific metric**
 - TE default metric

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—8-8

For the tunnel to dynamically discover its path through the network, the headend router must be provided with information on which to base this calculation. Specifically, it needs to be provided with the following:

- **Maximum bandwidth:** The maximum bandwidth is the amount of bandwidth that is available on each link in the network. Because there are priority levels for traffic tunnels, the availability information must be sent for each priority level for each link. Including priority levels means that the path decision mechanism is given the opportunity to choose a link with some bandwidth already allocated to a lower-priority tunnel, forcing that lower-priority tunnel to be “bounced” off the link.
- **Link resource class:** For administrative reasons, the network administrator may decide that some tunnels are not permitted to use certain links. To accomplish this goal, for each link, a link resource class must be defined and advertised. The definition of the tunnel may include a reference to particular “affinity bits.” The tunnel affinity bits are matched against the link resource class to determine whether a link may be used as part of the LSP.
- **Constraint-based specific metric:** Each link has a cost or metric for calculating routes in the normal operation of the IGP. It may be that, when calculating the LSP for traffic tunnels, the link should use a different metric. Thus, a constraint-based specific metric may be specified.

Constraint-Based Path Computation

This topic introduces constraint-based path computation.

Constraint-Based Path Computation

- **Constraint-based routing is demand-driven.**
- **Resource-reservation-aware routing paradigm:**
 - **Based on criteria including, but not limited to, network topology**
 - **Calculated at the edge of a network:**
 - **Modified Dijkstra's algorithm at tunnel headend (CSPF [Constraint-based SPF] or PCALC [path calculation]).**
 - **Output is a sequence of IP interface addresses (next-hop routers) between tunnel endpoints.**

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2-5-9

In traditional networks, the IGP calculates paths through the network based on the network topology alone. Routing is destination-based, and all traffic to a given destination from a given source uses the same path through the network. That path is based simply on what the IGP regards as the “least cost” between the two points (source and destination).

“Constraint-based routing” (CBR) is the term that is used most often for this approach. In some situations it is also referred to as a “constraint-based shortest path first” (CSPF) calculation or a “path calculation” (PCALC).

CBR behaves in these ways:

- It augments the use of link cost by also considering other factors, such as bandwidth availability or link attributes, when choosing the path to a destination.
- It tends to be carried out at the edge of the network, discovering a path across the core to some destination elsewhere at the other edge of the network. Typically, this discovery uses the CSPF calculation (a version of shortest path first [SPF] that is used by IS-IS and OSPF, but considering other factors besides cost, such as bandwidth availability).
- It produces a sequence of IP addresses that correspond to the routers that are used as the path to the destination; these addresses are the next-hop addresses for each stage of the path.

A consequence of CBR is that, from one source to one destination, many different paths can be used through the network, depending on the requirements of those data flows.

Constraint-Based Path Computation (Cont.)

- **Constraint-based routing takes into account:**
 - Policy constraints associated with the tunnel and physical links
 - Physical resource availability
 - Network topology state
- **Two types of tunnels can be established across those links with matching attributes:**
 - **Dynamic**—Using the least-cost path computed by OSPF or IS-IS
 - **Explicit**—Definition of a path by using Cisco IOS configuration commands

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2--8-10

When choosing paths through the network, the CBR system takes into account the following factors:


- The topology of the network, including information about the state of the links (the same information that is used by normal hop-by-hop routing)
- The resources that are available in the network, such as the bandwidth not already allocated on each link and at each of the eight priority levels (priority 0 to 7)
- The requirements that are placed on the constraint-based calculation that is defining the policy or the characteristics of this traffic tunnel

Of course, CBR is a dynamic process, which responds to a request to create a path and calculates (or recalculates) the path based on the status of the network at that time. The network administrator can also explicitly define the traffic tunnel.

By using commands such as **exclude-address** or **next-hop loose** in the explicit path configuration, the network administrator can mix static and dynamic computation.

Constraint-Based Path Computation (Cont.)

What is the best path from R1 to R6 with bandwidth of 30 Mbps?




An example network is shown in the figure. Each link specifies a link cost for metric calculation and a bandwidth available for reservation; for example, a metric of 10 and an available bandwidth of 100 Mbps is shown for the link between R1 and R2. Other than these criteria, no links are subject to any policy restriction that would disallow their use for creating traffic tunnels.

The requirement is to create a tunnel from R1 to R6 with a bandwidth of 30 Mbps.

Based simply on the link costs, the least-cost path from R1 to R6 is R1-R4-R6 with a cost of 30. However, the link from R4 to R6 has only 20 Mbps of bandwidth available for reservation and therefore cannot fulfill the requirements of the tunnel.

Similarly, the link R5-R6 has only 20 Mbps available as well, so no paths can be allocated via R5.

Constraint-Based Path Computation (Cont.)



The diagram now shows only those links that can satisfy the requirement for 30 Mbps of available bandwidth.

Over this topology, two tunnel paths are shown:

- The red (solid arrow) path shows the result of a dynamic constraint-based path calculation. The calculation has ignored any links that do not satisfy the bandwidth requirement (those shown in the previous figure but not shown here, such as the connections to R5) and then executes a CSPF calculation on what remains. This calculation has yielded the path R1-R2-R5-R6 with a path cost of 40.
- The network administrator has statically defined the blue (dashed arrow) path (R1-R4-R5-R6). Had the administrator attempted to define a path that did not have the required free bandwidth, tunnel establishment would have failed. This tunnel does indeed fulfill the minimum bandwidth requirement. However, adding the link costs yields a total of 45, which is not the lowest cost possible.

TE Processes

This topic explains the processes used in traffic engineering (TE).

Traffic Engineering Processes

- **Information distribution**
- **Path selection and calculation**
- **Path setup**
- **Trunk admission control**
- **Forwarding traffic on to tunnel**
- **Path maintenance**

© 2006 Cisco Systems, Inc. All rights reserved. MPLS v2.2—8-13

There are six TE processes to understand:

- **Information distribution:** Because the resource attributes are configured locally for each link, they must be distributed to the headend routers of traffic trunks. These resource attributes are flooded throughout the network using extensions to link-state intradomain routing protocols, either IS-IS and OSPF. The flooding takes place under these conditions:
 - Link-state changes occur.
 - The resource class of a link changes (this could happen when a network administrator reconfigures resource class of a link).
 - The amount of available bandwidth crosses one of the preconfigured thresholds.

The frequency of flooding is bounded by the OSPF and IS-IS timers. There are up thresholds and down thresholds. The up thresholds are used when a new trunk is admitted. The down thresholds are used when an existing trunk goes away.

Note Extensions to IS-IS and OSPF are discussed in a later topic.

- **Path selection:** Path selection for a traffic trunk takes place at the headend routers of traffic trunks. Using extended IS-IS or OSPF, the edge routers have knowledge of both network topology and link resources. For each traffic trunk, the router starts from the destination of the trunk and attempts to find the shortest path toward the source (that is, using the SPF algorithm). The SPF calculation does not consider the links that are explicitly excluded by the resource class affinities of the trunk or the links that have insufficient bandwidth. The output of the path selection process is an explicit route consisting of a sequence of label switching routers. This path is used as the input to the path setup procedure.
- **Path setup:** Path setup is initiated by the headend routers. Resource Reservation Protocol (RSVP) is the protocol that establishes the forwarding state along the path computed in the path selection process. The headend router sends a PATH message for each traffic trunk it originates.
- **Trunk admission control:** Trunk admission control handles the situation when a router along a computed path has insufficient bandwidth to honor the resource requested in the PATH message.
- **Forwarding of traffic to a tunnel:** Traffic can be forwarded to a tunnel by several means, including these:
 - Static routing
 - Policy routing from the global routing table
 - Autoroute
- **Path maintenance:** Path maintenance refers to two operations: path reoptimization and restoration.

Role of RSVP in Path Setup and Trunk Admission Control

This topic explains how RSVP is used in setting up an LSP path and in trunk admission control.

Role of RSVP in Path Setup Procedures

- **When the path has been determined, a signaling protocol is needed:**
 - To establish and maintain label-switched paths (LSPs) for traffic tunnels
 - For creating and maintaining resource reservation states across a network (bandwidth allocation)
- **The Resource Reservation Protocol (RSVP) was adopted by the MPLS workgroup of the IETF.**

© 2006 Cisco Systems, Inc. All rights reserved. MPLS v2.2—8-14

The result of the constraint-based calculation is a list of routers that form the path to the destination. The path is a list of IP addresses that identify each next hop along the path.

However, this list of routers is known only to the router at the headend of the tunnel that is attempting to build the tunnel. Somehow, this now-explicit path must be communicated to the intermediate routers. It is not up to the intermediate routers to make their own CSPF calculations; they merely abide by the path that is provided to them by the headend router.

Therefore, some signaling protocol is required to confirm the path, to check and apply the bandwidth reservations, and finally to apply the MPLS labels to form the MPLS LSP through the routers. The MPLS working group of the Internet Engineering Task Force (IETF) has adopted RSVP to confirm and reserve the path and apply the labels that identify the tunnel. Label Distribution Protocol (LDP) is used to apply the labels for the underlying MPLS network.

Note RSVP is needed for both explicit and dynamic path setup.

Path Setup with RSVP

RSVP is used for path setup.

Path Setup with RSVP

- **When the path has been calculated, it must be signaled across the network.**
 - Reserve any bandwidth to avoid “double booking” from other TE reservations.
 - Priority can be used to preempt low priority existing tunnels.
- **RSVP is used to set up TE LSP.**
 - PATH message (from head to tail) carries LABEL_REQUEST.
 - RESV message (from tail to head) carries LABEL.
- **When RESV messages reaches headend, tunnel interface is up.**
- **RSVP messages exist for LSP teardown and error signaling.**

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2--8-15

To signal the calculated path across the network, a PATH message is sent to the tailend router by the headend router for each traffic tunnel the headend originates.

Note This process occurs in the MPLS control plane.

The PATH message carries the explicit route (the output of the path selection process) computed for this traffic trunk, consisting of a sequence of label switching routers. The PATH message always follows this explicit route. Each intermediate router along the path performs trunk admission control after receiving the PATH message. When the router at the end of the path (tailend router) receives the PATH message, it sends a RESV message in the reverse direction toward the headend of the traffic trunk. As the RESV message flows toward the sender, each intermediate node reserves bandwidth and allocates labels for the trunk. When the RESV message reaches the sender, the LSP is already established.

RSVP messages also support for LSP teardown and error signaling.

Trunk Admission Control with RSVP

RSVP is also used in trunk admission control.

RSVP and Trunk Admission Control

- **On receipt of PATH message:**
 - Router checks whether there is bandwidth available to honor the reservation.
 - If bandwidth is available, then RSVP is accepted.
- **On receipt of a RESV message:**
 - Router actually reserves the bandwidth for the TE LSP.
 - If preemption is required, lower priority LSPs are torn down.
- **OSPF or IS-IS updates are triggered.**

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—8-16

Trunk admission control is used to confirm that each device along the computed path has sufficient provisioned bandwidth to support the resource requested in the PATH message. When a router receives a PATH message, it checks whether there is enough bandwidth to honor the reservation at the setup priority of the traffic trunk. Priority levels 0 to 7 are supported. If there is enough provisioned bandwidth, the reservation is accepted, otherwise the path setup fails. When the router receives the RESV message, it reserves bandwidth for the LSP. If preemption is required, the router must tear down existing trunks with a lower priority. As part of trunk admission control, the router must do local accounting to keep track of resource utilization and trigger IS-IS or OSPF updates when the available resource crosses the configured thresholds.

Forwarding Traffic to a Tunnel

This topic discusses changes that occur in the forwarding table when MPLS TE is implemented.

Forwarding Traffic to a Tunnel

- **IP routing is separate from LSP routing and does not see internal details of the LSP.**
- **The traffic has to be mapped to the tunnel:**
 - **Static routing**—The static route in the IP routing table points to an LSP tunnel interface.
 - **Policy routing**—The next-hop interface is an LSP tunnel.
 - **Autoroute**—SPF enhancement:
 - The headend sees the tunnel as a directly connected interface (for modified SPF only).
 - The default cost of a tunnel is equal to the shortest IGP metric regardless of the used path.

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2--9-17

The tunnel normally does not appear in the IP routing table. The IP routing process does not see the tunnel, so the tunnel is normally not included in any SPF calculations. The IP traffic can be mapped onto a tunnel in four ways:

- Using static routes that point to the tunnel interfaces.
- Using policy-based routing (PBR) and setting the next hop for the destination to the tunnel interface.
- Using the autoroute feature, an SPF enhancement that includes the tunnel interface in the route calculation as well. The result of the autoroute feature is that the tunnel is seen at the headend (and only there) as a directly connected interface. The metric (cost) of the tunnel is set to the normal IGP metric from the tunnel headend to the tunnel endpoint (over the least-cost path, regardless of whether the tunnel is actually using the least-cost path).

Note With the autoroute feature, the traffic-engineered tunnel appears in the IP routing table as well, but this appearance is restricted to the tunnel headend only.

- Using forwarding adjacency, which allows the tunnel to be announced via OSPF or IS-IS like any other unidirectional link (UDL). To be used for data forwarding, such a tunnel has to be set up bidirectionally.

The first two options are not very flexible or scalable. The traffic for each destination that needs to use the tunnel must be manually mapped to the tunnel.

For example, when you are using static routes, the tunnel is used only for the explicit static routes. Any other traffic that is not covered by the explicit static routes, including traffic for the tailend router (even though the tunnel terminates on that router), will not be able to use the tunnel; instead, it will follow the normal IGP path.

Forwarding Traffic to a Tunnel: Autoroute

This topic describes how the IP forwarding database is modified when the autoroute feature is implemented.

IP Forwarding Database Modification with Autoroute

- **Autoroute feature enables the headend to see the LSP as a directly connected interface:**
 - This feature is used only for the SPF route determination, not for the constraint-based path computation.
 - All traffic directed to prefixes topologically behind the tunnel endpoint (tailend) is forwarded onto the tunnel.
- **Autoroute affects the headend only; other routers on the LSP path do not see the tunnel.**
- **Tunnel is treated as a directly connected link to the tailend:**
 - When tunnel tail is seen in PATH list during IGP SPF, replace outgoing physical interface with tunnel interface.
 - Inherit tunnel to all downstream neighbors of tailend.

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—8-18

To overcome the problems that result from static routing configurations onto MPLS TE tunnels, the autoroute feature of Cisco IOS software was introduced. The autoroute feature enables the headend router to see the MPLS TE tunnel as a directly connected interface. The headend uses the MPLS TE in its modified SPF computations.


Note The MPLS TE tunnel is used only for normal IGP route calculation (at the headend only) and is not included in any constraint-based path computation.

When the tunnel is built, there is a directly connected link from headend to tailend.

The autoroute feature enables all the prefixes that are topologically behind the MPLS TE tunnel endpoint (tailend) to be reachable via the tunnel itself. This contrasts with static routing, where only statically configured destinations are reachable via the tunnel.

The autoroute feature affects the headend router only and has no effect on intermediate routers. These routers still use normal IGP routing for all the destinations.


Autoroute Topology (OSPF and ISIS)



The figure shows an example with two TE tunnels from R1. When the tunnels are up, R4 and R5 appear as directly connected neighbors to R1.

Note The tunnels are seen for routing purposes only by R1, the headend router. Intermediate routers do not see the tunnel nor do they take it into consideration for route calculations.

Autoroute Topology (OSPF and ISIS)



Summary

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

Summary

- **Traffic tunnels are configured with a set of resource requirements, such as bandwidth and priority.**
- **CSPF augments the link cost by considering other factors such as bandwidth availability or link latency when choosing a path.**
- **RSVP with TE extensions is used for establishing and maintaining LSPs.**
- **TE tunnels do not appear in the IP routing table.**

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—8-21

References

For additional information, refer to this resource:

- RFC 2746, *RSVP Operation over IP Tunnels*

Configuring MPLS TE on Cisco IOS Platforms

Overview

This lesson describes the Multiprotocol Label Switching (MPLS) traffic engineering (TE) commands for the implementation of MPLS traffic tunnels. The configuration commands that are needed to support MPLS TE are explained, and sample setups are presented.

Objectives

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

- Identify the tasks that are required to implement MPLS TE
- Enable device-level MPLS TE support
- Enable MPLS TE support in an IS-IS environment
- Enable MPLS TE support in an OSPF environment
- Enable MPLS TE on an interface
- Create and configure a traffic tunnel
- Enable traffic tunnels with autoroute

MPLS TE Configuration Road Map

This topic describes the requirements to implement MPLS TE on a Cisco IOS device.

MPLS TE Configuration Road Map

- **Use device supporting MPLS.**
 - **Implement ip cef global configuration command.**
- **Use IGP supporting MPLS TE.**
 - **OSPF**
 - **IS-IS**
- **Enable MPLS TE tunnels.**
 - **Configure mpls traffic-eng tunnels (global configuration mode).**
 - **Configure mpls traffic-eng tunnels (interface configuration mode).**
- **Create and configure traffic tunnels.**
- **Deploy traffic tunnels with autoroute.**

© 2006 Cisco Systems, Inc. All rights reserved. MPLS v2.2—8.3

MPLS has to be supported to implement MPLS TE, so Cisco Express Forwarding (CEF) is required.

MPLS TE has to be supported in the Interior Gateway Protocol (IGP). Open Shortest Path First (OSPF) and Intermediate System-to-Intermediate System (IS-IS) currently have MPLS TE support available. The configuration process for both IGPs is covered later in this lesson.

MPLS TE needs to be enabled both in global configuration mode and on specific interfaces. After MPLS TE is enabled, you need to create and configure traffic tunnels and then send the appropriate traffic into the tunnels.

Enabling Device-Level MPLS TE Support

This topic describes the command syntax that is required to enable device-level MPLS TE support on a Cisco IOS device.

Enabling Device-Level MPLS TE Support

```
router(config)#  
ip cef
```

- Enables IP CEF switching

```
router(config)#  
mpls traffic-eng tunnels
```

- Enables the MPLS TE tunnel feature on a device

© 2006 Cisco Systems, Inc. All rights reserved. MPLS v2.2—8-4

ip cef

To enable CEF on the router, use the **ip cef** command in global configuration mode. To disable CEF, use the **no** form of this command.

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

Syntax Description

Parameter	Description
distributed	(Optional) Enables distributed CEF (dCEF) operation. Distributes CEF information to line cards. Line cards perform express forwarding.

Usage Guidelines

CEF is an advanced Layer 3 IP switching technology. CEF optimizes network performance and scalability for networks with dynamic, topologically dispersed traffic patterns, such as those associated with web-based applications and interactive sessions.

Note	CEF is a prerequisite for MPLS traffic engineering.


mpls traffic-eng tunnels (global)

To enable MPLS TE tunnel signaling on a device, use the **mpls traffic-eng tunnels** command in global configuration mode. To disable this feature, use the **no** form of this command.

- **mpls traffic-eng tunnels**
- **no mpls traffic-eng tunnels**

Usage Guidelines

Enables the MPLS TE feature on a device. To use the feature, MPLS TE must also be enabled on the desired interfaces.



This figure illustrates the global configuration commands needed to enable device-level MPLS TE support.

Note These basic commands need to be configured on all provider edge (PE) and provider (P) routers in the MPLS network.

Enabling MPLS TE Support in IS-IS

This topic describes the command syntax that is required to enable MPLS TE support in an IS-IS implementation.

Enabling MPLS TE Support in IS-IS

```
router(config-router)#  
mpls traffic-eng {level-1 | level-2}
```

- Turns on MPLS traffic engineering for IS-IS Level 1 or Level 2

```
router(config-router)#  
mpls traffic-eng router-id {interface}
```

- Specifies the interface to be associated with the TE router identifier (also endpoint of a tunnel)

```
router(config-router)#  
metric-style wide
```

- Configures a router to generate and accept only new-style TLV objects

© 2006 Cisco Systems, Inc. All rights reserved.MPLS v2.2—8-8

mpls traffic-eng

To configure a router running IS-IS so that it floods MPLS TE link information into the indicated IS-IS level, use the **mpls traffic-eng** command in router configuration mode. To disable the flooding of MPLS TE link information into the indicated IS-IS level, use the **no** form of this command.

- **mpls traffic-eng** {level-1 | level-2}
- **no mpls traffic-eng** {level-1 | level-2}

Syntax Description

Parameter	Description
<i>level-1</i>	Floods MPLS TE link information into IS-IS Level 1
<i>level-2</i>	Floods MPLS TE link information into IS-IS Level 2

Usage Guidelines

This command appears as part of the routing protocol tree and causes link resource information (for instance, the bandwidth available) for appropriately configured links to be flooded in the IS-IS link-state database.

mpls traffic-eng router-id

To specify the TE router identifier that is to be the IP address that is associated with the given interface for the node, use the **mpls traffic-eng router-id** command in router configuration mode. To disable this feature, use the **no** form of this command.

- **mpls traffic-eng router-id** *interface*
- **no mpls traffic-eng router-id** *interface*

Syntax Description

Parameter	Description
<i>interface</i>	The MPLS TE router identifier is taken from the IP address of the supplied interface. This MPLS TE router identifier should be configured as the tunnel destination for tunnels that originate at other routers and terminate at this router. This interface should be a stable interface that will not go up and down, such as a loopback interface.

Usage Guidelines

The router identifier acts as a stable IP address for the TE configuration. This stable IP address is flooded to all nodes. For all TE tunnels that originate at other nodes and end at this node, the tunnel destination must be set to the TE router identifier of the destination node, because that identifier is the address that the TE topology database at the tunnel head uses for its path calculation.

Note	The MPLS router ID has to be the same address as the IGP router ID. Typically a loopback interface is used for the router ID.

It is a best practice to use the same loopback interface for MPLS TE, IGP, Border Gateway Protocol (BGP), and Label Distribution Protocol (LDP) router identification.

metric-style wide

To configure a router running IS-IS so that it generates and accepts only new-style type, length, and value (TLV) objects, use the **metric-style wide** command in router configuration mode. To disable this function, use the **no** form of this command.

- **metric-style wide** [**transition**] [**level-1** | **level-2** | **level-1-2**]
- **no metric-style wide** [**transition**] [**level-1** | **level-2** | **level-1-2**]

Syntax Description


Parameter	Description
transition	(Optional) Instructs the router to accept both old- and new-style TLV objects
level-1	Enables this command on routing Level 1
level-2	Enables this command on routing Level 2
level-1-2	Enables this command on routing Levels 1 and 2

Usage Guidelines

If you enter the **metric-style wide** command, a router generates and accepts only new-style TLV objects. The router uses less memory and other resources than it would if it were generating both old-style and new-style TLV objects.

This style is appropriate for enabling MPLS TE across an entire network.

Note This discussion of metric styles and transition strategies is oriented toward TE deployment. Other commands and models may be appropriate if the new-style TLV objects are desired for other reasons. For example, a network may require wider metrics but may not use TE.



This figure illustrates typical the router mode configuration commands needed to enable MPLS TE support in an IS-IS implementation.

Note These commands need to be configured on all PE and P routers in the MPLS network for an IS-IS implementation.

Enabling MPLS TE Support in OSPF

This topic describes the command syntax that is required to enable MPLS TE support in an OSPF implementation.

Enabling MPLS TE Support in OSPF

```
router(config-router)#  
mpls traffic-eng area area-num
```

- Turns on OSPF traffic engineering for the area

```
router(config-router)#  
mpls traffic-eng router-id interface
```

- Specifies the interface to be associated with the TE router identifier (also endpoint of a tunnel)

© 2006 Cisco Systems, Inc. All rights reserved. MPLS v2.2—5.8

mpls traffic-eng area

To configure a router running OSPF MPLS so that it floods TE for the indicated OSPF area, use the **mpls traffic-eng area** command in router configuration mode. To disable flooding of TE for the indicated OSPF area, use the **no** form of this command.

- **mpls traffic-eng area** *area-number*
- **no mpls traffic-eng area** *area-number*

Syntax Description

Parameter	Description
<i>area-number</i>	Indicates the OSPF area on which MPLS TE is enabled

Usage Guidelines

This command affects the operation of MPLS TE only if MPLS TE is enabled for that routing protocol instance. Currently, only a single area may be enabled for TE.

mpls traffic-eng router-id

To specify the TE router identifier for the node that is to be the IP address that is associated with the given interface, use the **mpls traffic-eng router-id** command in router configuration mode.

■ mpls traffic-eng router-id *interface*

Syntax Description

Parameter	Description
<i>interface</i>	The MPLS TE router identifier is taken from the IP address of the supplied interface. This MPLS TE router identifier should be configured as the tunnel destination for tunnels that originate at other routers and terminate at this router. This interface should be a stable interface that will not go up and down, such as a loopback interface.


Usage Guidelines

This router identifier acts as a stable IP address for the TE configuration. This stable IP address is flooded to all nodes. For all TE tunnels that originate at other nodes and end at this node, the tunnel destination must be set to the TE router identifier of the destination node, because that identifier is the address that the TE topology database at the tunnel head uses for its path calculation.

Note	The MPLS router ID has to be the same address as the IGP router ID. Typically a loopback interface is used for the router ID.

It is a best practice to use the same loopback interface for MPLS TE, IGP, BGP, and LDP router identification.

Implementing MPLS TE Support in OSPF



© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2-8-9

This figure illustrates the typical router mode configuration commands needed to enable MPLS TE support in an OSPF implementation.

Note These commands need to be configured on all PE and P routers in the MPLS network for an OSPF implementation.

Enabling Basic MPLS TE on an Interface

This topic describes the command syntax that is required to enable basic MPLS TE on an interface on a Cisco IOS device.

Enabling Basic MPLS TE on an Interface

```
router(config-if)#  
mpls ip
```

- Enables MPLS on an interface

```
router(config-if)#  
mpls traffic-eng tunnels
```

- Enables the MPLS TE tunnel feature on an interface

```
router(config-if)#  
ip rsvp bandwidth [interface-kbps [single-flow-kbps]]
```

- Enables RSVP for IP on an interface and specify amount of bandwidth to be reserved

© 2006 Cisco Systems, Inc. All rights reserved. MPLS v2.2—8-10

mpls ip

To enable MPLS forwarding of IP version 4 (IPv4) packets along normally routed paths for a particular interface, use the **mpls ip** command in interface configuration mode. To disable this feature, use the **no** form of this command.

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

mpls traffic-eng tunnels (interface)

To enable MPLS TE tunnel signaling on an interface, use the **mpls traffic-eng tunnels** command in interface configuration mode. To disable this feature, use the **no** form of this command.

- **mpls traffic-eng tunnels**
- **no mpls traffic-eng tunnels**

Note MPLS TE tunnel signaling must first be enabled at the global level.

Usage Guidelines

This command enables the MPLS TE feature on the interface. To use the feature, MPLS TE must also be enabled on the device globally. An enabled interface has its resource information flooded into the appropriate IGP link-state database and accepts TE tunnel signaling requests.

ip rsvp bandwidth

To enable Resource Reservation Protocol (RSVP) for IP on an interface, use the **ip rsvp bandwidth** interface configuration command. To disable RSVP, use the **no** form of this command.

- **ip rsvp bandwidth** [interface-kbps [single-flow-kbps]]
- **no ip rsvp bandwidth** [interface-kbps [single-flow-kbps]]

Syntax Description

Parameter	Description
<i>interface-kbps</i>	(Optional) Maximum amount of bandwidth, in kilobits per second, that may be allocated by RSVP flows. The range is from 1 to 10,000,000.
<i>single-flow-kbps</i>	(Optional) Maximum amount of bandwidth, in kilobits per second, that may be allocated to a single flow. The range is from 1 to 10,000,000.


Usage Guidelines

RSVP is disabled by default. If the **ip rsvp bandwidth** command is entered but no bandwidth values are supplied (for example, if you enter **ip rsvp bandwidth** and then press the Enter key), a default bandwidth value is assumed for both the *interface-kbps* and *single-flow-kbps* arguments.

RSVP cannot be configured with Versatile Interface Processor (VIP) dCEF.

Note	To make MPLS TE work over a subinterface, you must configure RSVP on the main interface with the ip rsvp bandwidth command even if the main interface is not used for MPLS TE.

Implementing Basic MPLS TE on an Interface



This figure illustrates the commands to enable basic MPLS TE support on an interface.

Note These commands need to be configured on all interfaces between the PE and P routers in the MPLS network.

Creating and Configuring a Traffic Tunnel

This topic describes the command syntax that is required to create and configure a traffic tunnel on a Cisco IOS device.

Creating a Tunnel

```
router(config)#  
interface tunnel number
```

- Configures a tunnel interface

```
router(config-if)#  
ip unnumbered type number
```

- Enables IP on the interface, and brings it up without a unique IP address

```
router(config-if)#  
tunnel destination ip-address
```

- Specifies the destination for the tunnel
- The tunnel source command is not needed.

© 2006 Cisco Systems, Inc. All rights reserved. MPLS v2.2—8-12

interface tunnel

The **interface tunnel** command is used to declare a tunnel interface. The tunnel interface *number* parameter value is in the range of 0 to 2147483647.

- **interface tunnel** *number*

ip unnumbered

To enable IP processing on an interface without assigning an explicit IP address to the interface, use the **ip unnumbered** command in interface configuration mode or subinterface configuration mode. To disable IP processing on the interface, use the **no** form of this command.

- **ip unnumbered** *type number*
- **no ip unnumbered** *type number*

Note Although the tunnel interface could be assigned an IP address with the **ip address** command, using an unnumbered interface is considered a best practice.

Syntax Description

Parameter	Description
<i>type</i>	Type of another interface on which the router has an assigned IP address. The interface cannot be another unnumbered interface.
<i>number</i>	Number of another interface on which the router has an assigned IP address. The interface cannot be another unnumbered interface.

Usage Guidelines

Whenever the unnumbered interface generates a packet (for example, for a routing update), it uses the address of the specified interface as the source address of the IP packet. It also uses the address of the specified interface in determining which routing processes are sending updates over the unnumbered interface.

For MPLS TE tunnels, it is recommended that you use a loopback interface for addressing the unnumbered interface.

tunnel destination

To specify the destination for a tunnel interface, use the **tunnel destination** interface configuration command.

- **tunnel destination** {*hostname* | *ip-address*}

Syntax Description

Parameter	Description
<i>hostname</i>	Name of the host destination
<i>ip-address</i>	IP address of the host destination expressed in decimal in four-part dotted notation

Usage Guidelines


For MPLS TE, it is a good practice to use the router ID of the remote router as the destination address.

Note	Two tunnels cannot use the same encapsulation mode with exactly the same source and destination address.

With MPLS traffic engineering, the source for the tunnel is set automatically.

Note	The tunnel source command is not used.

Example: Creating a Tunnel Interface



© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—8-13

This figure illustrates the commands to create a tunnel interface.

Note These commands are configured only on the ingress PE for the MPLS TE tunnel.

Configuring a Traffic Tunnel

This topic describes the command syntax that is required to create and configure a traffic tunnel on a Cisco IOS device.

Creating and Configuring a Traffic Tunnel

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

```
tunnel mode mpls traffic-eng
```

- Sets the tunnel encapsulation mode to MPLS TE

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

```
tunnel mpls traffic-eng bandwidth bandwidth
```

- Configures the requested bandwidth for the MPLS TE tunnel

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

```
tunnel mpls traffic-eng priority setup-priority [hold-priority]
```

- Configures the setup and reservation priority for an MPLS TE tunnel

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—8-14

tunnel mode mpls traffic-eng

To set the mode of a tunnel to MPLS for TE, use the **tunnel mode mpls traffic-eng** command in interface configuration mode.

- **tunnel mode mpls traffic-eng**

Usage Guidelines

This command specifies that the tunnel interface is for an MPLS TE tunnel and enables the various tunnel MPLS configuration options.

tunnel mpls traffic-eng bandwidth

To configure the bandwidth that is required for an MPLS TE tunnel, use the **tunnel mpls traffic-eng bandwidth** command in the interface configuration mode. To disable this feature, use the **no** form of this command.

- **tunnel mpls traffic-eng bandwidth** [**sub-pool** | **global**] *bandwidth*
- **no tunnel mpls traffic-eng bandwidth** [**sub-pool** | **global**] *bandwidth*

Syntax Description

Parameter	Description
<i>bandwidth</i>	The bandwidth required for an MPLS TE tunnel. Bandwidth is specified in kilobits per second. Bandwidth, in kilobits per second, is set aside for the MPLS traffic engineering tunnel. The range is between 1 and 4294967295.
sub-pool	(Optional) Indicates a subpool tunnel.
global	(Optional) Indicates a global pool tunnel. Entering this keyword is not necessary, for all tunnels are global pool in the absence of the sub-pool keyword. But if users of pre-DiffServ-aware TE images enter this keyword, it is accepted.

Usage Guidelines

This command specifies that the tunnel bandwidth for an MPLS TE tunnel.

Default

The default bandwidth is 0.

tunnel mpls traffic-eng priority

To configure the setup and reservation priority for an MPLS TE tunnel, use the **tunnel mpls traffic-eng priority** command in interface configuration mode. To remove the specified setup and reservation priority, use the **no** form of this command.

- **tunnel mpls traffic-eng priority** *setup-priority* [*hold-priority*]
- **no tunnel mpls traffic-eng priority** *setup-priority* [*hold-priority*]

Syntax Description

Parameter	Description
<i>setup-priority</i>	The priority that is used when signaling a label-switched path (LSP) for this tunnel to determine which existing tunnels are eligible for preemption. The range is from 0 to 7, where a lower numeric value indicates a higher priority.
<i>hold-priority</i>	(Optional) The priority that is associated with an LSP for this tunnel when it has been established to determine whether it should be preempted by other LSPs that are being signaled. The range is from 0 to 7, where a lower numeric value indicates a higher priority.

Usage Guidelines

When an LSP is being signaled and an interface does not currently have enough bandwidth available for that LSP, the call admission software preempts lower-priority LSPs so that the new LSP can be admitted. (LSPs are preempted if doing so allows the new LSP to be admitted.)

In the described determination, the priority of the new LSP is its setup priority, and the priority of the existing LSP is its hold priority. The two priorities make it possible to signal an LSP with a low setup priority (so that the LSP does not preempt other LSPs on setup) but a high hold priority (so that the LSP is not preempted after it is established).

Setup priority and hold priority are typically configured to be equal, and setup priority cannot be better (numerically smaller) than the hold priority.

Defaults

The defaults are as follows:

- Setup priority = 7
- Hold priority = 7

Creating and Configuring a Traffic Tunnel (Cont.)

```
router(config)#
```

```
ip explicit-path {name word | identifier number}  
[{enable | disable }]
```

- Enters the subcommand mode for IP explicit paths and creates or modifies the specified path

```
router(cfg-ip-expl-path)#
```

```
next-address ip-address
```

- Specifies the next IP address in the explicit path

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

```
tunnel mpls traffic-eng path-option number {dynamic |  
explicit {name path-name | path-number}} [lockdown]
```

- Configures the tunnel to use a named IP explicit path or a path dynamically calculated from the TE topology database

© 2006 Cisco Systems, Inc. All rights reserved.MPLS v2.2—8-15

ip explicit-path

To enter the subcommand mode for IP explicit paths to create or modify the named path, use the **ip explicit-path** command in global configuration mode. An IP explicit path is a list of IP addresses, each representing a node or link in the explicit path.

- **ip explicit-path** {*name word* | **identifier number**} [{**enable** | **disable**}]

Syntax Description

Parameter	Description
name <i>word</i>	Specifies an explicit path by name.
identifier <i>number</i>	Specifies an explicit path by number. You can specify a number from 1 to 65535.
enable	(Optional) Sets the state of the path to be enabled.
disable	(Optional) Prevents the path from being used for routing while it is being configured.

next-address

To specify the next IP address in the explicit path, use the **next-address** command in IP explicit path configuration mode. To remove the specified next IP address in the explicit path, use the **no** form of this command.

- **next-address** *ip-address*
- **no next-address** *ip-address*

Default

The next IP address in the explicit path is not specified.

tunnel mpls traffic-eng path-option

To configure a path option for an MPLS-TE tunnel, use the **tunnel mpls traffic-eng path-option** command in interface configuration mode. To disable the specified path option, use the **no** form of this command.

- **tunnel mpls traffic-eng path-option** *number* {**dynamic** | **explicit** {**name** *path-name* | *path-number*}} [**lockdown**]
- **no tunnel mpls traffic-eng path-option** *number* {**dynamic** | **explicit** {**name** *path-name* | *path-number*}} [**lockdown**]

Syntax Description

Parameter	Description
<i>number</i>	When multiple path options are configured, lower-numbered options are preferred.
dynamic	Path of the LSP is dynamically calculated.
name <i>path-name</i>	Path name of the IP explicit path that the tunnel uses with this option.
<i>path-number</i>	Path number of the IP explicit path that the tunnel uses with this option.
lockdown	(Optional) The LSP cannot be reoptimized.

Usage Guidelines

Multiple path setup options may be configured for a single tunnel. For example, you can configure several explicit paths and a dynamic option for one tunnel. Path setup prefers options with lower numbers to options with higher numbers, so option 1 is the most strongly preferred option. The explicit path is configured using the **ip explicit-path** command.

Mapping Traffic into Tunnels with Autoroute

This section identifies the command syntax that is required to deploy traffic tunnels with the autoroute feature on Cisco IOS devices.

Deploying Traffic Tunnels with Autoroute

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

```
tunnel mpls traffic-eng autoroute announce
```

- **Causes the IGP to use the tunnel in its enhanced SPF**

© 2006 Cisco Systems, Inc. All rights reserved.MPLS v2.2—8-18

tunnel mpls traffic-eng autoroute announce

To instruct the IGP to use the tunnel in its shortest path first (SPF) calculation if the tunnel is up, use the **tunnel mpls traffic-eng autoroute announce** command in interface configuration mode. To disable this feature, use the **no** form of this command.

- **tunnel mpls traffic-eng autoroute announce**
- **no tunnel mpls traffic-eng autoroute announce**

Default


The tunnel is not used by the IGP in its SPF next-hop calculation unless a mechanism to map traffic to the tunnel is implemented.

Usage Guidelines

The **tunnel mpls traffic-eng autoroute announce** command is a scalable way to map IP traffic onto a tunnel. Alternate mapping methods include static routes pointing to the tunnel interface and using policy-based routing (PBR) to set the next hop for the destination to the tunnel interface.

For flexibility, the autoroute feature is preferred over these alternatives.

Example: Configuring TE Options on a Tunnel Interface



```
ip explicit-path name via_P1 enable
next-address 10.1.1.12
next-address 10.1.1.2

interface tunnel 0
ip unnumbered Loopback0
tunnel destination 10.1.1.2
tunnel mode mpls traffic-eng
tunnel mpls traffic-eng priority 3 3
tunnel mpls traffic-eng bandwidth 128
tunnel mpls traffic-eng path-option 1 explicit name via_P1
tunnel mpls traffic-eng autoroute announce
```

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—8-17

This figure illustrates commands to configure some TE options on a tunnel interface.

Note These commands are configured only on the ingress PE for the MPLS TE tunnel.

Summary

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

Summary

- **Required tasks to implement TE in a MPLS VPN network include:**
 - Enabling device-level MPLS TE support on PE and P routers
 - Enabling MPLS TE support in an IS-IS environment on PE and P routers
 - Enabling MPLS TE support in an OSPF environment on PE and P routers
 - Enabling MPLS TE on an interface on PE and P routers
 - Creating and configuring a traffic tunnel on ingress PE router
 - Enabling traffic tunnels with autoroute in ingress PE router

© 2006 Cisco Systems, Inc. All rights reserved. MPLS v2.2--8-18

References

For additional information, refer to this resource:

- MPLS Configuration Examples and TechNotes.
http://www.cisco.com/en/US/tech/tk436/tk428/tech_configuration_examples_list.html

Monitoring Basic MPLS TE on Cisco IOS Platforms

Overview

This lesson presents some of the commands, syntax, and descriptions for monitoring Multiprotocol Label Switching (MPLS) traffic engineering (TE). This lesson explains how to monitor simple TE tunnels in an MPLS network to ensure that traffic engineering is functioning appropriately.

Objectives

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

- Monitor MPLS TE tunnels
- Monitor MPLS TE, including verifying RSVP information on an interface, verifying that MPLS TE is present in the global routing table, and verifying that IP traffic is forwarded through the MPLS TE tunnels

Monitoring MPLS TE Tunnels

This topic describes how to monitor MPLS TE tunnels.

Monitoring MPLS TE Tunnels

```
router#  
show ip rsvp interface
```

- Confirms which interfaces have RSVP information configured

```
router#  
show mpls traffic-eng tunnels tunnel_interface [brief]
```

- Shows summary information about a tunnel

```
router#  
show mpls traffic-eng tunnels [parameters]
```

- Shows detailed information about a tunnel

© 2006 Cisco Systems, Inc. All rights reserved. MPLS v2.2—7.3

show ip rsvp interface

To display Resource Reservation Protocol (RSVP)-related information, use the **show ip rsvp interface** command in EXEC mode.

- **show ip rsvp interface** [*interface-type interface-number*] [**detail**]

Syntax Description

Parameter	Description
<i>interface-type</i>	(Optional) Type of interface
<i>interface-number</i>	(Optional) Number of interface
detail	(Optional) Displays additional information about interfaces

The example shows output from the **show ip rsvp interface** command.

Monitoring MPLS TE Tunnels: show ip rsvp interface

```
P72#show ip rsvp interface
interface    allocated    i/f max    flow max    sub max
Se0/0        0            128K       128K        0
Se0/0.111    0            128K       128K        0
Se0/0.112    0            128K       128K        0
Se0/0.272    0            128K       128K        0
```

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—7-4

show mpls traffic-eng tunnels

To show information about MPLS TE tunnels, use the **show mpls traffic-eng tunnels** command.

- **show mpls traffic-eng tunnels** *tunnel-interface* [**brief**]
- **show mpls traffic-eng tunnels** [**destination** *address*]
[**source-id** {*number* | *ip-address* | *ip-address number*}]
[**role** {**all** | **head** | **middle** | **tail** | **remote**}] [**up** | **down**] [**name** *string*]
[**suboptimal constraints** {**none** | **current** | **max**}] [[**interface in** *physical-interface*]
[**interface out** *physical-interface*] | [**interface** *physical-interface*] [**brief**]

The table describes the parameters available for the **show mpls traffic-eng tunnels** command.

Syntax Description

Parameter	Description
<i>tunnel-interface</i>	Displays information for the specified tunneling interface.
brief	(Optional) Displays the information in brief format.
destination <i>address</i>	(Optional) Restricts the display to tunnels that are destined to the specified IP address.
source-id	(Optional) Restricts the display to tunnels with a matching source IP address or tunnel number.
<i>number</i>	(Optional) Tunnel number.
<i>ip-address</i>	(Optional) Source IP address.
<i>ip-address number</i>	(Optional) Source IP address and tunnel number.
role	(Optional) Restricts the display to tunnels with the indicated role (all , head , middle , tail , or remote).
all	(Optional) Displays all tunnels.
head	(Optional) Displays tunnels with their heads at this router.
middle	(Optional) Displays tunnels with their midpoints at this router.
tail	(Optional) Displays tunnels with their tails at this router.
remote	(Optional) Displays tunnels with their heads at another router; this display is a combination of the middle and tail keyword values.
up	(Optional) Displays tunnels if the tunnel interface is up. Tunnel midpoints and tails are typically up or not present.
down	(Optional) Displays tunnels that are down.
name <i>string</i>	(Optional) Displays tunnels with the specified name. The tunnel name is derived from the interface description, if specified; otherwise, it is the interface name. The tunnel name is included in the signaling message so that it is available at all hops.
suboptimal constraints none	(Optional) Displays tunnels whose path metric is greater than the shortest unconstrained path. Selected tunnels have a longer path than the shortest path of the Interior Gateway Protocol (IGP).
suboptimal constraints current	(Optional) Displays tunnels whose path metric is greater than the current shortest path, constrained by the configured tunnel options. Selected tunnels would have a shorter path if they were reoptimized immediately.
suboptimal constraints max	(Optional) Displays tunnels whose path metric is greater than the current shortest path, constrained by the configured tunnel options and considering only the capacity of the network. Selected tunnels would have a shorter path if no other tunnels were consuming network resources.
interface in <i>physical-interface</i>	(Optional) Displays tunnels that use the specified input interface.
interface out <i>physical-interface</i>	(Optional) Displays tunnels that use the specified output interface.
interface <i>physical-interface</i>	(Optional) Displays tunnels that use the specified interface as an input or output interface.
brief	(Optional) Specifies one line per tunnel.

The example shows output from the **show mpls traffic-eng tunnels brief** command.

Monitoring MPLS TE Tunnels: show mpls traffic-eng tunnels brief

```
PE71#show mpls traffic-eng tunnels brief
Signaling Summary:
  LSP Tunnels Process:      running
  RSVP Process:            running
  Forwarding:              enabled
  Periodic reoptimization:  every 30 seconds, next in 11 seconds
  Periodic auto-bw collection: disabled

TUNNEL NAME          DESTINATION   UP IF      DOWN IF    STATE/PROT
PE71_t0              192.168.7.33  -          Se0/0.11   up/up
PE72_t0              192.168.7.17  Se0/0.11   -          up/up
Displayed 1 (of 1) heads, 0 (of 0) midpoints, 1 (of 1) tails
```

Note Depending on its length, the tunnel name may be truncated.

Monitoring MPLS TE

This topic describes how to monitor MPLS TE.

Monitoring MPLS TE

```
router#  
show mpls traffic-eng autoroute
```

- Shows tunnels that are announced to IGP, including interface, destination, and bandwidth

```
router#  
show ip cef network [mask]
```

- Shows network entries in the Forwarding Information Base (FIB) to verify that IP traffic is forwarded through tunnel interface

```
router#  
show ip cef vrf vrf-name network [mask]
```

- Shows VRF network entries in the FIB to verify that MPLS VPN traffic is forwarded through tunnel interface

© 2006 Cisco Systems, Inc. All rights reserved. MPLS v2.2—7-6

show mpls traffic-eng autoroute

To show tunnels that are announced to IGP, including interface, destination, and bandwidth, use the **show mpls traffic-eng autoroute** command in privileged EXEC mode.

- **show mpls traffic-eng autoroute**

Usage Guidelines

The enhanced shortest path first (SPF) calculation of the IGP has been modified so that it uses TE tunnels. This command shows which tunnels the IGP is currently using in its SPF, or which tunnels that are up and have autoroute configured.

The example shows output from the **show mpls traffic-eng autoroute** command.

Monitoring MPLS TE: show mpls traffic-eng autoroute

```
PE71#show mpls traffic-eng autoroute
MPLS TE autorouting enabled
destination 0000.0000.0172.00 has 1 tunnels
  Tunnel0 (load balancing metric 20000000, nexthop 192.168.7.33)
```

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—7-7

Note The list of tunnels is organized by destination. All tunnels to a destination will carry a share of the traffic to that destination.

show ip cef *network*

To verify that IP traffic is forwarded through a tunnel interface, use the **show ip cef *network*** command in privileged EXEC mode.

- **show ip cef *network* [*mask*]**

The example shows output from the **show ip cef *network*** command.

Monitoring MPLS TE: show ip cef *network*

```
PE71#show ip cef 192.168.7.33
192.168.7.33/32, version 25, epoch 0
0 packets, 0 bytes
tag information set, shared
  local tag: 30
  fast tag rewrite with Tu0, point2point, tags imposed: {33}
via 192.168.7.33, Tunnel0, 11 dependencies
  next hop 192.168.7.33, Tunnel0
  valid adjacency
  tag rewrite with Tu0, point2point, tags imposed: {33}
```

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—7-8

show ip cef vrf *vrf-name* network

To verify that Virtual Private Network (VPN) traffic is forwarded through a tunnel interface, use the **show ip cef vrf *vrf-name* network** command in privileged EXEC mode.

■ **show ip cef vrf *vrf-name* network [mask]**

The example shows output from the **show ip cef vrf *vrf-name* network** command.

Monitoring MPLS TE: show ip cef vrf *vrf-name* network

```
PE71#show ip cef vrf CustA 10.1.72.49
10.1.72.49/32, version 20, epoch 0
0 packets, 0 bytes
tag information set
  local tag: VPN-route-head
  fast tag rewrite with Tu0, point2point, tags imposed: {33 34}
via 192.168.7.33, 0 dependencies, recursive
  next hop 192.168.7.33, Tunnel0 via 192.168.7.33/32
  valid adjacency
  tag rewrite with Tu0, point2point, tags imposed: {33 34}
```

© 2006 Cisco Systems, Inc. All rights reserved.

MPLS v2.2—7-9

Summary

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

Summary

- **Use the following commands to monitor MPLS TE tunnels:**
 - show ip rsvp interface
 - show mpls traffic-eng tunnels
- **Use the following commands to monitor MPLS TE:**
 - show mpls traffic-eng autoroute
 - show ip cef *network*
 - show ip cef vrf *vrf-name network*

© 2006 Cisco Systems, Inc. All rights reserved. MPLS v2.2—7-10

References

For additional information, refer to this resource:

- MPLS Configuration Examples and TechNotes.
http://www.cisco.com/en/US/tech/tk436/tk428/tech_configuration_examples_list.html

Module Summary

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

Module Summary

- **Traffic engineering measures, models, and controls traffic**
- **The MPLS TE uses traffic tunnels, RSVP with TE extensions, and a link state IGP**
- **MPLS TE implementation includes device, interface, routing protocol, and tunnel configuration**
- **Commands including show ip cef and show mpls can be used to monitor MPLS TE**

© 2006 Cisco Systems, Inc. All rights reserved. MPLS v2.2—7-1

This module provides a brief introduction to Multiprotocol Label Switching (MPLS) traffic engineering (TE).

References

For additional information, refer to this resource:

- MPLS Configuration Examples and TechNotes.
http://www.cisco.com/en/US/tech/tk436/tk428/tech_configuration_examples_list.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) What is the main motivation for implementing traffic engineering? (Source: Introducing the TE Concept)
-
- Q2) Which two situations can result in network congestion? (Choose two.) (Source: Introducing the TE Concept)
- A) The least-cost network path is being used to route traffic.
 - B) Traffic streams are efficiently mapped onto available resources.
 - C) Network resources are insufficient to accommodate the offered load.
 - D) Traffic streams are inefficiently mapped onto available resources.
- Q3) Short traffic bursts are best handled by which three mechanisms? (Choose three.) (Source: Introducing the TE Concept)
- A) rate limiting
 - B) traffic engineering
 - C) queuing algorithms
 - D) intelligent packet dropping
- Q4) When you are using TE with a Layer 2 overlay model, which two of the following transport traffic across a network? (Choose two.) (Source: Introducing the TE Concept)
- A) DVCs
 - B) PVCs
 - C) RVCs
 - D) SVCs
- Q5) When you are using a traffic-engineered Layer 3 model, which two of the following are limitations? (Choose two.) (Source: Introducing the TE Concept)
- A) using static routes
 - B) using policy routing
 - C) unavailability of explicit routing
 - D) path computation that is based just on the IGP metric
- Q6) When you are using MPLS-TE, a minimum of _____ labels is assigned to a packet at the ingress router. (Source: Introducing the TE Concept)
- A) one
 - B) two
 - C) three
 - D) four

- Q7) A set of data flows that share some common feature, attribute, or requirement is called _____. (Source: Understanding MPLS TE Components)
- A) a static route
 - B) a policy route
 - C) a TE tunnel
 - D) TE LSP
- Q8) The constraint-based path computation uses which algorithm? (Source: Understanding MPLS TE Components)
- A) DUAL algorithm
 - B) modified Dijkstra algorithm
 - C) modified Bell-Howell algorithm
 - D) none of the above
- Q9) Which two can be used to advertise a traffic tunnel so that it will appear in the IP routing table? (Choose two.) (Source: Understanding MPLS TE Components)
- A) autoroute feature
 - B) TE mapping statement
 - C) static routes
 - D) mpls traffic-eng tunnels command
- Q10) What is the role of RSVP in an MPLS TE implementation? (Source: Understanding MPLS TE Components)
- A) It identifies the best path for the tunnel.
 - B) It reserves the bandwidth required by the tunnel.
 - C) It performs the CBR calculations for the tunnel setup.
 - D) It assigns the label for the MPLS LSP.
- Q11) To enable MPLS TE tunnel signaling on a device, you must use the _____ command. (Source: Configuring MPLS TE on Cisco IOS Platforms)
- A) **ip cef**
 - B) **mpls ip**
 - C) **mpls-te enable**
 - D) **mpls traffic-eng tunnels**
- Q12) The command to allow a router to generate and accept only new-style TLV objects under IS-IS is the _____ command. (Source: Configuring MPLS TE on Cisco IOS Platforms)
- A) **ip cef**
 - B) **metric-style wide**
 - C) **mpls traffic-eng**
 - D) **mpls traffic-eng area**
- Q13) Which command enables MPLS TE in an OSPF implementation? (Source: Configuring MPLS TE on Cisco IOS Platforms)
- A) **mpls traffic-eng tunnels**
 - B) **metric-style wide**
 - C) **mpls traffic-eng area**
 - D) **mpls traffic-eng**

- Q14) The _____ command is used to declare a TSP tunnel interface. (Source: Configuring MPLS TE on Cisco IOS Platforms)
- A) interface tunnel *number*
 - B) **tunnel destination** *ip-address*
 - C) **tunnel mode mpls traffic-eng**
 - D) **tunnel mpls traffic-eng autoroute announce**
- Q15) The _____ command is used to instruct the IGP to use the tunnel in its SPF or next-hop calculation. (Source: Configuring MPLS TE on Cisco IOS Platforms)
- A) interface tunnel *number*
 - B) **tunnel destination** *ip-address*
 - C) **tunnel mode mpls traffic-eng**
 - D) **tunnel mpls traffic-eng autoroute announce**
- Q16) The _____ command is used to display Resource Reservation Protocol (RSVP)-related information. (Source: Monitoring Basic MPLS TE on Cisco IOS Platforms)
- A) show mpls traffic-eng tunnels
 - B) show ip rsvp interface
 - C) show ip cef network
 - D) show ip rsvp network
- Q17) The _____ command is used to verify that IP traffic is forwarded through a tunnel interface. (Source: Monitoring Basic MPLS TE on Cisco IOS Platforms)
- A) show mpls traffic-eng tunnels
 - B) show ip rsvp interface
 - C) show ip cef *network*
 - D) show ip rsvp *network*
- Q18) The _____ command is used to show information about MPLS TE tunnels. (Source: Monitoring Basic MPLS TE on Cisco IOS Platforms)
- A) show mpls traffic-eng tunnels
 - B) show ip rsvp interface
 - C) show ip cef network
 - D) show ip rsvp network

Answer Key

- Q1) Cost reduction
- Q2) C, D
- Q3) A, C, D
- Q4) B, D
- Q5) C, D
- Q6) B
- Q7) C
- Q8) B
- Q9) A, C
- Q10) C
- Q11) D
- Q12) B
- Q13) C
- Q14) A
- Q15) D
- Q16) B
- Q17) C
- Q18) A

