

# CCNP Implementing Secured Converged WANs (ISCW 642-825) Lab Portfolio

**David Kotfila**

**Joshua Moorhouse**

**Ross Wolfson, CCIE No. 16696**

**Cisco Press**

800 East 96th Street

Indianapolis, Indiana 46240 USA

# CCNP Implementing Secured Converged WANs (ISCW 642-825) Lab Portfolio

David Kotfila ▪ Joshua Moorhouse ▪ Ross Wolfson

Copyright© 2008 Cisco Systems, Inc.

Published by:

Cisco Press

800 East 96th Street

Indianapolis, IN 46240 USA

All rights reserved. No part of this book may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or by any information storage and retrieval system, without written permission from the publisher, except for the inclusion of brief quotations in a review.

Printed in the United States of America

First Printing March 2008

Library of Congress Cataloging-in-Publication Data

Kotfila, David A.

CCNP implementing secured converged WANs (ISCW 642-825) lab portfolio / David Kotfila, Joshua Moorhouse, Ross Wolfson.

p. cm.

ISBN 978-1-58713-215-5 (pbk.)

1. Extranets (Computer networks) 2. Wide area networks (Computer networks) I. Moorhouse, Joshua D. II. Wolfson, Ross G. III. Title.

TK5105.875.E87K64 2008

005.8—dc22

2008006745

ISBN-13: 978-158713-215-5

ISBN-10: 1-58713-215-X

## Warning and Disclaimer

This book provides labs consistent with the Cisco Networking Academy CCNP Implementing Secured Converged WANs (ISCW 642-825) curriculum. Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied.

The information is provided on an “as is” basis. The authors, Cisco Press, and Cisco Systems, Inc., shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this book or from the use of the discs or programs that may accompany it.

The opinions expressed in this book belong to the authors and are not necessarily those of Cisco Systems, Inc.

**Publisher**  
Paul Boger

**Associate Publisher**  
Dave Dusthimer

**Cisco Representative**  
Anthony Wolfenden

**Cisco Press Program Manager**  
Jeff Brady

**Executive Editor**  
Mary Beth Ray

**Managing Editor**  
Patrick Kanouse

**Senior Development Editor**  
Christopher Cleveland

**Senior Project Editor**  
Tonya Simpson

**Copy Editor**  
Sheri Cain

**Technical Editors**  
Tommy Crawford,  
Geovany González

**Editorial Assistant**  
Vanessa Evans

**Book Designer**  
Louisa Adair

**Proofreader**  
Paula Lowell

## Trademark Acknowledgments

All terms mentioned in this book that are known to be trademarks or service marks have been appropriately capitalized. Cisco Press or Cisco Systems, Inc., cannot attest to the accuracy of this information. Use of a term in this book should not be regarded as affecting the validity of any trademark or service mark.

## Corporate and Government Sales

The publisher offers excellent discounts on this book when ordered in quantity for bulk purchases or special sales, which may include electronic versions and/or custom covers and content particular to your business, training goals, marketing focus, and branding interests. For more information, please contact: **U.S. Corporate and Government Sales** 1-800-382-3419  
corpsales@pearsontechgroup.com

For sales outside the United States, please contact: **International Sales**  
international@pearsoned.com

## Feedback Information

At Cisco Press, our goal is to create in-depth technical books of the highest quality and value. Each book is crafted with care and precision, undergoing rigorous development that involves the unique expertise of members from the professional technical community.

Readers' feedback is a natural continuation of this process. If you have any comments regarding how we could improve the quality of this book, or otherwise alter it to better suit your needs, you can contact us through e-mail at [feedback@ciscopress.com](mailto:feedback@ciscopress.com). Please make sure to include the book title and ISBN in your message.

We greatly appreciate your assistance.

This book is part of the Cisco Networking Academy® series from Cisco Press. The products in this series support and complement the Cisco Networking Academy curriculum. If you are using this book outside the Networking Academy, then you are not preparing with a Cisco trained and authorized Networking Academy provider.

For more information on the Cisco Networking Academy or to locate a Networking Academy, please visit [www.cisco.com/edu](http://www.cisco.com/edu).



**Americas Headquarters**  
Cisco Systems, Inc.  
170 West Tasman Drive  
San Jose, CA 95134-1706  
USA  
[www.cisco.com](http://www.cisco.com)  
Tel: 408 526-4000  
800 553-NETS (6387)  
Fax: 408 527-0883

**Asia Pacific Headquarters**  
Cisco Systems, Inc.  
168 Robinson Road  
#28-01 Capital Tower  
Singapore 068912  
[www.cisco.com](http://www.cisco.com)  
Tel: +65 6317 7777  
Fax: +65 6317 7799

**Europe Headquarters**  
Cisco Systems International BV  
Haarlerbergpark  
Haarlerbergweg 13-19  
1101 CH Amsterdam  
The Netherlands  
[www-europe.cisco.com](http://www-europe.cisco.com)  
Tel: +31 0 800 020 0791  
Fax: +31 0 20 357 1100

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at [www.cisco.com/go/offices](http://www.cisco.com/go/offices).

©2007 Cisco Systems, Inc. All rights reserved. CCVP, the Cisco logo, and the Cisco Square Bridge logo are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn is a service mark of Cisco Systems, Inc.; and Access Registrar, Aironet, BPX, Catalyst, CCDA, CCDP, CCIE, CCIIP, CCNA, CCNP, CCSP, 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, Follow Me Browsing, FormShare, GigaDrive, GigaStack, HomeLink, Internet Quotient, IOS, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, iQuick Study, LightStream, Linksys, MeetingPlace, MGX, Networking Academy, Network Registrar, Packet, PIX, ProConnect, RateMUX, ScriptShare, SlideCast, SMARTnet, StackWise, 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. (0609R)

## Introduction

My first motivation for writing this book was to serve the needs of CCNP instructors and students in the Cisco Networking Academy Program. For the past four years, I (David) have had the privilege of serving on the National Advisory Council for the Cisco Networking Academy, representing four-year colleges and universities. Also on the council are numerous two-year community colleges. Inevitably, at council meetings, we discussed both CCNP curriculum and labs. As I spoke with a number of my CCNP instructor peers, a common theme emerged. Instructors felt that the labs needed to be rewritten to be more comprehensive. In the past, labs have lacked complexity. When I realized that I was rewriting the Networking Academy CCNP labs, and that my peers were rewriting the same labs, the thought occurred to me that perhaps an engineering school, like RPI, was up to the task of writing these labs in a way that would better serve the needs of the community. It is not that the previous labs were inappropriate; rather, it's just that the Cisco Networking Academy has grown up. Having just celebrated its tenth birthday, folks in the Academy are ready for bigger challenges. I believe that these labs fill that role.

My second motivation for writing these labs was to help networking professionals who are trying to upgrade their skill set to the CCNP level. As a former hiring manager at a Tier 1 ISP, I have a strong sense of what an industry is looking for when it hires someone with CCNP credentials. Each year, numerous hiring managers from Fortune 500 companies contact me about hiring my students. I know the level of expertise they expect from a CCNP. These labs reflect the convictions those managers shared with me.

My third motivation for writing these labs was to see how much of a challenge university undergraduates could rise to if they were asked to do a big job. My coauthors, Josh Moorhouse and Ross Wolfson, were both undergraduates when they authored these labs. I gave them a huge task, and they responded with skill and grace. I firmly believe that we frequently do not ask enough of our students. If we ask for greatness, we will sometimes get it. If we settle for the normal, we are more assured of success, but we might miss the opportunity to see our students soar to new heights. With these labs, whether you are an instructor or student, I hope that your technical knowledge soars to new heights.

## Goals and Methods

The most important goal of this book is to help you master the technologies necessary to configure secure WANs in a production environment. After all, what is the point of getting certified and getting that dream job or promotion if you cannot perform after you are there? Although it is impossible to simulate a network of 300 routers, we have added loopback interfaces to simulate additional networks and increase complexity.

This book's secondary goal is to help people pass the ISCW certification exam. For two years, I was on the CCNP Assessment authoring team. After all of those years of complaining, "What were they thinking when they put *that* question on the exam?," suddenly, the questions I was writing were the subject of someone else's complaint. I know how important it is, both to students and networking professionals, to pass certifications. Frequently, prestige, promotion, and money are all at stake. Although all the core configurations on the certification exam are covered in this book, no static document, like a book, can keep up with the dynamic way in which the certification exam is constantly upgraded.

---

## Who Should Read This Book?

Cisco Networking Academy instructors and students who want a written copy of the electronic labs will find this book greatly useful. In addition to all the official labs that are part of the Networking Academy curriculum, additional Challenge and Troubleshooting labs have been added to test your mastery.

Networking professionals, either in formal classes or studying alone, will also find great value in this book. Knowing how expensive it can be to purchase your own lab equipment, as many labs as possible were written with only three routers. (To adequately cover some topics, four routers were necessary.) Final configurations were included with each lab so that even if you do not have all the equipment, you can walk through the configurations in your head.

## What You Need to Configure the Labs

These labs were written on four Cisco 2811 routers using the IOS image `c2800nm-advipservicesk9-mz.124-10.bin`.

You should be able to configure the labs on any Cisco router that uses a 12.4 advanced IP services image of the IOS.

Classes and individuals using older Cisco devices (or less robust versions of the IOS) might find that some of the commands are different or not supported.

Example: It is not possible to run the 12.4 release of the advanced IP services IOS image on a Cisco 2600 Series router. It is possible to run this image on a Cisco 2600XM router if you upgrade the Flash and RAM and can obtain the new IOS image.

## How This Book Is Organized

People preparing for the ISCW certification exam should work through this book cover-to-cover. Networking professionals needing help or a refresher on a particular topic can skip right to the area in which they need assistance.

The chapters cover the following topics:

- **Chapter 1, “Remote Network Connectivity Requirements”**—This chapter covers design concepts associated with remote-network connectivity. No labs are associated with it. However, there is a walk-through of the lab setup that is used throughout this book.
- **Chapter 2, “Teleworker Connectivity”**—The equipment necessary for configuring Point-to-Point Protocol over Ethernet (PPPoE) is physically different than the hardware necessary for the labs in the rest of this curriculum. Networking Academy students can simulate configuring this equipment using a Flash application, `dsl_standalone`. (Networking professionals can use a sample configuration if they do not have access to this application.)
- **Chapter 3, “IPsec VPNs”**—Cisco, in recognition of how difficult it is to stay current in all the protocols that a network engineer needs to stay current in, advocates the use of their GUI configuration tool, Security Device Manager (SDM). Like any GUI tool that creates configurations for you, it’s easy to use. Also, like any GUI tool, times arise when the GUI produces unexpected configurations and/or side effects. Therefore, it is necessary to know how to use the GUI (to save time) and how to edit the command-line interface (CLI) for the times when the GUI produces problematic results. This chapter’s labs teach you both skills.

- **Chapter 4, “Frame Mode MPLS Implementation”**—Multiprotocol Label Switching (MPLS) is a technology that is growing in its deployment. The basic lab for this chapter is possible using only three routers. MPLS virtual private networks (VPN) are also common. Although the certification requirement is only that you be able to describe (not configure) MPLS VPNs, we have included an optional lab on how MPLS VPNs are configured. Unfortunately, it is necessary to have five routers to really see what is occurring on the Internet service provider’s (ISP) side of the configuration. If you do not have this much hardware, you can still get a reasonable understanding of how MPLS VPNs work by merely reading this lab.
- **Chapter 5, “Cisco Device Hardening”**—When the first routers rolled off the production line, the burning issue wasn’t security. It was how we could get these devices to easily talk to each other using a variety of different protocols. Therefore, by default, many protocols and services were automatically turned on. As the Internet has matured, security has become a primary concern. Therefore, it is now necessary to turn off these services unless they are being used. Two tools help you accomplish this: One-Step Lockdown and AutoSecure. However, if you have a network situation that is somewhat unique—and who doesn’t?—you also need to know the CLI commands so that you can edit the generic configurations that these tools generate. These labs teach you both.
- **Chapter 6, “Cisco IOS Threat Defense Features”**—The labs demonstrate how to configure Cisco IOS Firewall and Intrusion Prevention System (IPS). As with previous chapters, you see how to configure them using SDM and how to configure and edit them by using CLI.
- **Chapter 7, “Case Studies”**—The first case study requires you to configure IPsec and Frame-Mode MPLS using CLI. The second case study requires you to configure Cisco IOS Firewall and IPS. As in previous chapters, you are asked to do some of this using SDM and do other tasks using CLI.

## NETLAB+ Compatibility

NDG has worked closely with the Cisco Networking Academy CCNP lab team to develop ISCW labs that are compatible with the installed base of NETLAB AE router pods. For current information on labs compatible with NETLAB+ go to <http://www.netdevgroup.com/ae/labs.htm>.

# Frame Mode MPLS Implementation

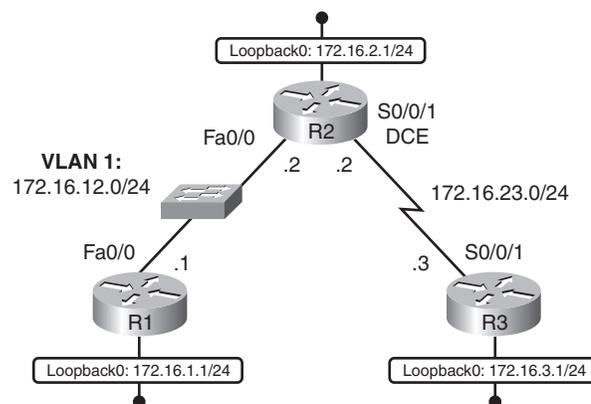
## Lab 4-1: Configuring Frame Mode MPLS (4.5.1)

In this lab, you learn how to do the following:

- Configure EIGRP on a router.
- Configure LDP on a router.
- Change the size of the MTU.
- Verify MPLS behavior.

Figure 4-1 illustrates the topology that is used for this lab.

**Figure 4-1** Topology Diagram



## Scenario

In this lab, you configure a simple Enhanced Interior Gateway Routing Protocol (EIGRP) network to route IP packets. You run Multiprotocol Label Switching (MPLS) over the IP internetwork to fast-switch Layer 2 frames.

## Step 1: Configure Addressing

Configure the loopback interfaces with the addresses shown in Figure 4-1. Also, configure the serial interfaces shown in the figure. Set the clock rate on the appropriate interface and issue the **no shutdown** command on all serial connections. Verify that you have connectivity across the local subnet by using the **ping** command:

```
R1(config)# interface loopback 0
R1(config-if)# ip address 172.16.1.1 255.255.255.0
R1(config-if)# interface fastethernet 0/0
R1(config-if)# ip address 172.16.12.1 255.255.255.0
R1(config-if)# no shutdown
-----
R2(config)# interface loopback 0
R2(config-if)# ip address 172.16.2.1 255.255.255.0
R2(config-if)# interface fastethernet 0/0
R2(config-if)# ip address 172.16.12.2 255.255.255.0
R2(config-if)# no shutdown
R2(config-if)# interface serial 0/0/1
R2(config-if)# ip address 172.16.23.2 255.255.255.0
R2(config-if)# clockrate 64000
R2(config-if)# no shutdown
-----
R3(config)# interface loopback 0
R3(config-if)# ip address 172.16.3.1 255.255.255.0
R3(config-if)# interface serial 0/0/1
R3(config-if)# ip address 172.16.23.3 255.255.255.0
R3(config-if)# no shutdown
```

## Step 2: Configure EIGRP AS 1

Configure EIGRP for AS1 on all three routers. Add the whole major network 172.16.0.0 and disable automatic summarization:

```
R1(config)# router eigrp 1
R1(config-router)# no auto-summary
R1(config-router)# network 172.16.0.0
-----
R2(config)# router eigrp 1
R2(config-router)# no auto-summary
R2(config-router)# network 172.16.0.0
-----
R3(config)# router eigrp 1
R3(config-router)# no auto-summary
R3(config-router)# network 172.16.0.0
```

EIGRP neighbor adjacencies should form between R1 and R2 and between R2 and R3. If the adjacencies do not form, troubleshoot by checking your interface configuration, EIGRP configuration, and physical connectivity.

What impact does IP connectivity have on MPLS?

---



---



---

### Step 3: Observe CEF Operation

Because all the routers have EIGRP adjacencies and are advertising the entire major 172.16.0.0 network, all routers should have full routing tables:

R1# **show ip route**

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route
```

Gateway of last resort is not set

```
172.16.0.0/24 is subnetted, 5 subnets
D    172.16.23.0 [90/2172416] via 172.16.12.2, 00:01:56, FastEthernet0/0
C    172.16.12.0 is directly connected, FastEthernet0/0
C    172.16.1.0 is directly connected, Loopback0
D    172.16.2.0 [90/156160] via 172.16.12.2, 00:01:56, FastEthernet0/0
D    172.16.3.0 [90/2300416] via 172.16.12.2, 00:01:51, FastEthernet0/0
```

On R1, if you perform a **tracert** to the R3's loopback, you see the path that the packet follows. This output changes slightly after you configure MPLS:

R1# **tracert 172.16.3.1**

Type escape sequence to abort.

Tracing the route to 172.16.3.1

```
 1 172.16.12.2 0 msec 0 msec 0 msec
 2 172.16.23.3 16 msec 12 msec *
```

Cisco Express Forwarding (CEF) is the Cisco proprietary Layer 3 switching algorithm for Cisco IOS routers. CEF allows forwarding to be distributed throughout the line cards on Cisco models, such as the Catalyst 6500. CEF also provides quicker switching than switching based on the routing table (process switching) or switching based on a standards-compliant forwarding information base (fast-switching).

What is the function of CEF?

---

---

---

Which information does CEF view as significant in making a forwarding determination for an IP packet?

---

You can also see that CEF is enabled by default by using the **show ip cef command**:

R1# **show ip cef**

Prefix	Next Hop	Interface
0.0.0.0/0	drop	Null0 (default route handler entry)
0.0.0.0/32	receive	
172.16.1.0/24	attached	Loopback0
172.16.1.0/32	receive	
172.16.1.1/32	receive	
172.16.1.255/32	receive	
172.16.2.0/24	172.16.12.2	FastEthernet0/0
172.16.3.0/24	172.16.12.2	FastEthernet0/0
172.16.12.0/24	attached	FastEthernet0/0
172.16.12.0/32	receive	
172.16.12.1/32	receive	
172.16.12.2/32	172.16.12.2	FastEthernet0/0
172.16.12.255/32	receive	
172.16.23.0/24	172.16.12.2	FastEthernet0/0
224.0.0.0/4	drop	
224.0.0.0/24	receive	
255.255.255.255/32	receive	

Another important CEF command is the **show ip cef non-recursive** command, which allows the user to display CEF forwarding information for prefixes installed in the routing table:

R1# **show ip cef non-recursive**

Prefix	Next Hop	Interface
172.16.1.0/24	attached	Loopback0
172.16.2.0/24	172.16.12.2	FastEthernet0/0
172.16.3.0/24	172.16.12.2	FastEthernet0/0
172.16.12.0/24	attached	FastEthernet0/0
172.16.12.2/32	172.16.12.2	FastEthernet0/0
172.16.23.0/24	172.16.12.2	FastEthernet0/0

CEF records both the Layer 3 next-hop information and the Layer 2 frame next-hop information. CEF currently supports the following Layer 2 protocols: Asynchronous Transfer Mode (ATM), Frame Relay, Ethernet, Fiber Distributed Data Interface (FDDI), Point-to-Point Protocol (PPP), High-Level Data Link Control (HDLC), and tunnels.

CEF is critical to the operation of MPLS on Cisco routers because MPLS packets must be forwarded based on label. Because the CEF architecture can support multiple protocols, such as IPv4 and IPv6, CEF switching could naturally be extended to support MPLS labels.

CEF should be enabled by default. If CEF is not enabled, issue the **ip cef** command in global configuration mode on each router.

## Step 4: Enable MPLS on All Physical Interfaces

MPLS is a standardized protocol that allows routers to switch packets based on labels, rather than route switch packets based on standards in the protocol's routing formula. Under normal IP routing, every intermediate system looks up the destination prefix of an IP packet in the Routing Information Base (RIB) of a router or in the Forwarding Information Base (FIB) of a fast switch at every Layer 3 node. Instead of switching that is based on prefix, the first router running MPLS can encapsulate the IP packet in an MPLS frame and then further encapsulate the packet in the Layer 2 frame before sending it across one of many supported Layer 2 media. At the next MPLS-enabled label switch router (LSR), the MPLS frame is read, and the IP packet is switched as an MPLS frame from router to router with little rewrite at each node.

This allows routers to switch multiple protocols (hence, the name) by using the same switching mechanism, as well as perform some other functionality not available in traditional destination-based forwarding, including Layer 2 VPNs (AToM), Layer 3 VPNs, and traffic engineering. MPLS runs between Layers 2 and 3 of the OSI model and, because of this, it is sometimes said to run at Layer 2 1/2. The MPLS header is 4 bytes long and includes a 20-bit label.

Configuring the interface-level command **mpls ip** on an interface tells the router to switch MPLS packets inbound and outbound on that interface and attempt to bring up MPLS adjacencies with the Label Distribution Protocol (LDP) out that egress interface. LDP facilitates communication between MPLS peers by allowing them to inform each other of labels to assign packets to particular destinations based on Layer 2, Layer 3, or other significant information.

Configure MPLS on all physical interfaces in the topology.

---

**Note:** If you run Cisco IOS Software Release 12.4 on your routers, use the **mpls ip** command in this lab. However, when Cisco first developed packet-labeling technology, it was called tag switching. Therefore, if you run an older version of IOS, you might see one of two different variations: The first variation is that your router accepts the **mpls ip** command. However, the commands are stored in IOS as **tag-switching** commands. The second variation is that your router won't accept the **mpls ip** command. In this event, the **mpls ip** command can be entered as the **tag-switching ip** command. Try the newer commands first, beginning with the **mpls** keyword.

---

```
R1(config)# interface fastethernet0/0
```

```
R1(config-if)# mpls ip
```

---

```
R2(config)# interface fastethernet0/0
```

```
R2(config-if)# mpls ip
```

```
*Jan 31 08:28:54.315: %LDP-5-NBRCHG: LDP Neighbor 172.16.1.1:0 (1) is UP
```

```
R2(config-if)# interface serial0/0/1
```

```
R2(config-if)# mpls ip
```

---

```
R3(config)# interface serial0/0/1
```

```
R3(config-if)# mpls ip
```

```
*Jan 31 08:32:11.571: %LDP-5-NBRCHG: LDP Neighbor 172.16.2.1:0 (1) is UP
```

Notice that as you configure MPLS on both ends of a connection, IOS logs a messages to the console on both routers, which indicates that an LDP neighbor adjacency has formed.

Although you are going to use LDP in this lab, another Cisco proprietary label-exchanging protocol exists: Tag Distribution Protocol (TDP), which was part of the Cisco tag-switching architecture. To change the protocol being used, use the **mpls label protocol protocol** command either on a global level at the global configuration prompt or on a per-interface basis, using the interface-level version of this command. Cisco TDP and MPLS LDP are nearly identical in functionality, but they use incompatible message formats and some different procedures. Cisco is changing from TDP to a fully compliant LDP.

## Step 5: Verify MPLS Configuration

MPLS has many **show** commands that you can use to verify proper MPLS operation. Issue the **show mpls interfaces** command to see a quick summary of interfaces configured with MPLS. Keep in mind that you see this output because you applied the **mpls ip** command to these interfaces:

R1# **show mpls interfaces**

Interface	IP	Tunnel	Operational
FastEthernet0/0	Yes (ldp)	No	Yes

R2# **show mpls interfaces**

Interface	IP	Tunnel	Operational
FastEthernet0/0	Yes (ldp)	No	Yes
Serial0/0/1	Yes (ldp)	No	Yes

R3# **show mpls interfaces**

Interface	IP	Tunnel	Operational
Serial0/0/1	Yes (ldp)	No	Yes

Issue the **show mpls ldp discovery** command to find out local sources for LDP exchanges and the **show mpls ldp neighbor** command to show LDP adjacencies. Notice that MPLS chooses its IDs based on loopback interfaces, similar to other protocols, such as Open Shortest Path First (OSPF) Protocol and Border Gateway Protocol (BGP):

R1# **show mpls ldp discovery**

Local LDP Identifier:

172.16.1.1:0

Discovery Sources:

Interfaces:

FastEthernet0/0 (ldp): xmit/rcv

LDP Id: 172.16.2.1:0; no host route

R1# **show mpls ldp neighbor**

Peer LDP Ident: 172.16.2.1:0; Local LDP Ident 172.16.1.1:0

TCP connection: 172.16.2.1.49525 - 172.16.1.1.646

State: Oper; Msgs sent/rcvd: 29/26; Downstream

Up time: 00:16:40

LDP discovery sources:

FastEthernet0/0, Src IP addr: 172.16.12.2

```
Addresses bound to peer LDP Ident:
```

```
172.16.12.2 172.16.23.2 172.16.2.1
```

---

```
R2# show mpls ldp discovery
```

```
Local LDP Identifier:
```

```
172.16.2.1:0
```

```
Discovery Sources:
```

```
Interfaces:
```

```
FastEthernet0/0 (ldp): xmit/rcv
```

```
LDP Id: 172.16.1.1:0; no host route
```

```
Serial0/0/1 (ldp): xmit/rcv
```

```
LDP Id: 172.16.3.1:0; no host route
```

```
R2# show mpls ldp neighbor
```

```
Peer LDP Ident: 172.16.1.1:0; Local LDP Ident 172.16.2.1:0
```

```
TCP connection: 172.16.1.1.646 - 172.16.2.1.49525
```

```
State: Oper; Msgs sent/rcvd: 27/30; Downstream
```

```
Up time: 00:17:06
```

```
LDP discovery sources:
```

```
FastEthernet0/0, Src IP addr: 172.16.12.1
```

```
Addresses bound to peer LDP Ident:
```

```
172.16.12.1 172.16.1.1
```

```
Peer LDP Ident: 172.16.3.1:0; Local LDP Ident 172.16.2.1:0
```

```
TCP connection: 172.16.3.1.34352 - 172.16.2.1.646
```

```
State: Oper; Msgs sent/rcvd: 27/26; Downstream
```

```
Up time: 00:16:23
```

```
LDP discovery sources:
```

```
Serial0/0/1, Src IP addr: 172.16.23.3
```

```
Addresses bound to peer LDP Ident:
```

```
172.16.3.1
```

---

```
R3# show mpls ldp discovery
```

```
Local LDP Identifier:
```

```
172.16.3.1:0
```

```
Discovery Sources:
```

```
Interfaces:
```

```
Serial0/0/1 (ldp): xmit/rcv
```

```
LDP Id: 172.16.2.1:0; no host route
```

```
R3# show mpls ldp neighbor
```

```
Peer LDP Ident: 172.16.2.1:0; Local LDP Ident 172.16.3.1:0
```

```
TCP connection: 172.16.2.1.646 - 172.16.3.1.34352
```

```
State: Oper; Msgs sent/rcvd: 27/28; Downstream
```

```
Up time: 00:17:19
```

```
LDP discovery sources:
```

```
Serial0/0/1, Src IP addr: 172.16.23.2
Addresses bound to peer LDP Ident:
172.16.12.2 172.16.23.2 172.16.2.1
```

What interface does LDP use on R1 to identify itself to other LDP peers?

---

---

---

What transport protocol does LDP use to communicate with other LDP peers?

---

---

In the configuration you set up in Step 4, all routers are acting as LSRs and running LDP. On LSRs, each forwarding equivalence class (in this case, each routable IP prefix) is assigned an MPLS label. LDP automatically distributes labels to peers to be used when sending traffic to specific destinations through the LSR. After the labels are distributed, switching for MPLS packets is done through the Label Information Base (LIB).

Display the contents of the LIB by using the **show mpls ldp bindings** command. A binding exists for every routed prefix; however, the bindings can vary from router to router because they can get swapped at each hop. In a larger network, the way labels are swapped is easier to see. The LIB is also referred to on Cisco routers as the TIB, which is a legacy name from tag switching. Do not be alarmed to see the LIB entries listed instead as TIB entries; this does not signal that TDP is the protocol being used for distribution:

```
R1# show mpls ldp bindings
tib entry: 172.16.1.0/24, rev 6
    local binding: tag: imp-null
    remote binding: tsr: 172.16.2.1:0, tag: 16
tib entry: 172.16.2.0/24, rev 8
    local binding: tag: 17
    remote binding: tsr: 172.16.2.1:0, tag: imp-null
tib entry: 172.16.3.0/24, rev 10
    local binding: tag: 18
    remote binding: tsr: 172.16.2.1:0, tag: 17
tib entry: 172.16.12.0/24, rev 4
    local binding: tag: imp-null
    remote binding: tsr: 172.16.2.1:0, tag: imp-null
tib entry: 172.16.23.0/24, rev 2
    local binding: tag: 16
    remote binding: tsr: 172.16.2.1:0, tag: imp-null
```

---

```
R2# show mpls ldp bindings
tib entry: 172.16.1.0/24, rev 6
    local binding: tag: 16
    remote binding: tsr: 172.16.1.1:0, tag: imp-null
    remote binding: tsr: 172.16.3.1:0, tag: 17
```

```

tib entry: 172.16.2.0/24, rev 8
    local binding: tag: imp-null
    remote binding: tsr: 172.16.1.1:0, tag: 17
    remote binding: tsr: 172.16.3.1:0, tag: 18
tib entry: 172.16.3.0/24, rev 10
    local binding: tag: 17
    remote binding: tsr: 172.16.1.1:0, tag: 18
    remote binding: tsr: 172.16.3.1:0, tag: imp-null
tib entry: 172.16.12.0/24, rev 4
    local binding: tag: imp-null
    remote binding: tsr: 172.16.1.1:0, tag: imp-null
    remote binding: tsr: 172.16.3.1:0, tag: 16
tib entry: 172.16.23.0/24, rev 2
    local binding: tag: imp-null
    remote binding: tsr: 172.16.1.1:0, tag: 16
    remote binding: tsr: 172.16.3.1:0, tag: imp-null

```

---

R3# **show mpls ldp bindings**

```

tib entry: 172.16.1.0/24, rev 6
    local binding: tag: 17
    remote binding: tsr: 172.16.2.1:0, tag: 16
tib entry: 172.16.2.0/24, rev 8
    local binding: tag: 18
    remote binding: tsr: 172.16.2.1:0, tag: imp-null
tib entry: 172.16.3.0/24, rev 10
    local binding: tag: imp-null
    remote binding: tsr: 172.16.2.1:0, tag: 17
tib entry: 172.16.12.0/24, rev 4
    local binding: tag: 16
    remote binding: tsr: 172.16.2.1:0, tag: imp-null
tib entry: 172.16.23.0/24, rev 2
    local binding: tag: imp-null
    remote binding: tsr: 172.16.2.1:0, tag: imp-null

```

The local bindings are generated by LDP on an LSR when LDP is enabled. A label is generated for every prefix in the routing table. These labels are then sent to all the router's LDP peers. A tag of implicit-NULL (imp-null in the output of the command **show mpls ldp bindings**) is advertised when the packet will not be forwarded locally based on label, but based on prefix. This situation regularly occurs with connected networks.

For example, assume R2 and R3 have already peered with each other using LDP. Now R1 begins running MPLS and attempts to peer to R2:

- R1 generates the locally bound label, namely 18, for the prefix 172.16.3.0/24 in its routing table.
- R1 advertises the local binding to its LDP peer, R2.

- R2 enters R1’s binding for the 172.16.3.0/24 prefix, now classified as a remote binding, into its LIB, regardless of whether it uses it to reach the destination network. The remote binding for this IP prefix through R1 is label 18.
- Based on the routing table, R2 will use R3 as the next hop for 172.16.3.0/24. R2 will not forward IP packets inside an MPLS encapsulation, but rather simply as IP packets because R3 has advertised the label of implicit-NULL to R2.

What is the significance of the local binding entry?

---



---

What is the significance of a remote binding entry?

---



---



---

On R2, why is there more than one remote binding for each of the networks in Figure 4-1?

---



---



---

Note that LDP assigns local labels to *all* Interior Gateway Protocol (IGP) prefixes and advertises the bindings to *all* LDP peers. The concept of split horizon does not exist; an LDP peer assigns its own local label to a prefix and advertises that back to the other LDP peer, even though that other LDP peer owns the prefix (it is a connected prefix) or that other LDP peer is the downstream LSR.

What is the meaning of the implicit-NULL label?

---



---



---



---

As previously mentioned, **tracert** would differ slightly after MPLS is set up. The output now includes labels for each hop. Unfortunately, because of the size of this network, you only see one label. In a larger network, you would see more hops and, therefore, more labels:

R1# **tracert 172.16.3.1**

Type escape sequence to abort.

Tracing the route to 172.16.3.1

```

 1 172.16.12.2 [MPLS: Label 17 Exp 0] 44 msec 44 msec 48 msec
 2 172.16.23.3 12 msec 12 msec *
```

## Step 6: Change MPLS MTU

Because you are adding extra header information to packets, the maximum transmission unit (MTU) of packets can change. Remember that each MPLS header is 4 bytes. The default MTU size of MPLS packets is taken from the interface it is running on, which, in the case of Ethernet, is 1500 bytes. To verify this, use the **show mpls interfaces interface-type interface-number detail** command to see the Ethernet connections of R1 and R2:

```
R1# show mpls interfaces fastethernet 0/0 detail
```

```
Interface FastEthernet0/0:
  IP labeling enabled (ldp):
    Interface config
  LSP Tunnel labeling not enabled
  BGP tagging not enabled
  Tagging operational
  Fast Switching Vectors:
    IP to MPLS Fast Switching Vector
    MPLS Turbo Vector
  MTU = 1500
```

```
R2# show mpls interfaces fastethernet 0/0 detail
```

```
Interface FastEthernet0/0:
  IP labeling enabled (ldp):
    Interface config
  LSP Tunnel labeling not enabled
  BGP tagging not enabled
  Tagging operational
  Fast Switching Vectors:
    IP to MPLS Fast Switching Vector
    MPLS Turbo Vector
  MTU = 1500
```

For this lab, we change the Ethernet connection between R1 and R2 to support two MPLS headers, so we will change the MPLS MTU to 1508 on their Fast Ethernet interfaces. To change the MPLS MTU, use the **mpls mtu size** command in interface configuration mode. Verify the change by using the **show mpls interfaces interface detail** command previously used:

```
R1(config)# interface fastethernet 0/0
```

```
R1(config-if)# mpls mtu 1508
```

```
R2(config)# interface fastethernet0/0
```

```
RR2(config-if)# mpls mtu 1508
```

```
R1# show mpls interface fastethernet 0/0 detail
```

```
Interface FastEthernet0/0:
  IP labeling enabled (ldp):
    Interface config
  LSP Tunnel labeling not enabled
  BGP tagging not enabled
  Tagging operational
```

```
Fast Switching Vectors:  
  IP to MPLS Fast Switching Vector  
  MPLS Turbo Vector  
MTU = 1508
```

---

```
R2# show mpls interface fastethernet 0/0 detail
```

```
Interface FastEthernet0/0:  
  IP labeling enabled (ldp):  
    Interface config  
  LSP Tunnel labeling not enabled  
  BGP tagging not enabled  
  Tagging operational  
  Fast Switching Vectors:  
    IP to MPLS Fast Switching Vector  
    MPLS Turbo Vector  
MTU = 1508
```

## Lab 4-2: Challenge Lab: Implementing MPLS VPNs (4.5.2)

This lab qualifies as a challenge lab rather than a required lab because the implementation of this lab is out of the scope of the CCNP2: Implementing Secure Converged WANs course. The requirement for the CCNP2 course is that individuals be able to *describe* MPLS VPN technology. There is no requirement to be able to *configure* it. The Border Gateway Protocol (BGP) commands used in this lab are beyond the scope of the commands learned in the CCNP1 course. Because MPLS VPN configuration typically takes place as part of an internal ISP network, this level of configuration knowledge is not required at the CCNP level.

Also, the lab requires five Layer 3 devices to show the distributed configuration and nature of MPLS. This is more equipment than is required in the Networking Academy bundle. There is nothing extraordinary about the configurations on the HQ and Branch routers; therefore, it is possible to solve this issue by using either Layer 3 switches or older routers. Only SP1, SP2, and SP3 need to be MPLS-capable routers. However, no required lab can exceed the standard equipment bundle.

This lab is included because, from a pedagogical perspective, it is often easier to understand difficult concepts through hands-on practice. If you choose not to actually configure devices, just reading through the configurations helps you understand and describe MPLS VPN technology.

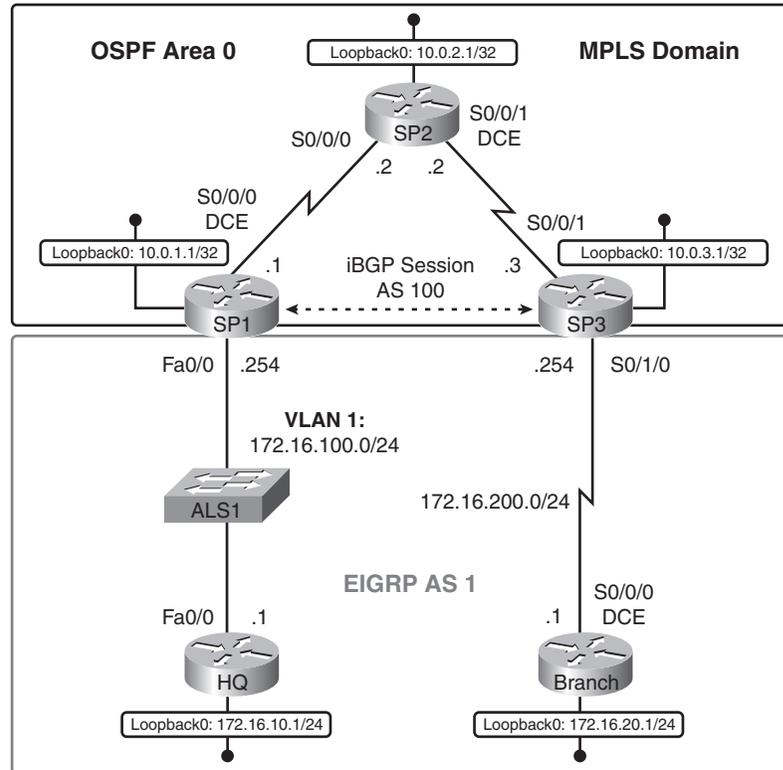
Also, some CCNPs are called upon to configure MPLS VPNs in the core of large enterprise networks or at ISPs. They will benefit from being exposed to this configuration in a lab environment.

In this lab, you learn how to do the following:

- Configure OSPF and EIGRP on a router.
- Enable MPLS on a router.
- Verify MPLS implementation.
- Configure a VRF instance.
- Use MBGP to exchange VPN routing updates.
- Verify VPN activity.

Figure 4-2 illustrates the topology that is used for this lab.

Figure 4-2 Topology Diagram



## Scenario

As a network engineer at a service-provider corporation, you suggest rolling out MPLS as a new transport technology to facilitate VPNs between customer sites that connect through your network. Your CIO has asked you to implement proof-of-concept in a lab environment, starting with a small implementation of MPLS VPNs before moving up to more moderately sized test cases.

MPLS VPN technology is a powerful technology that leverages the multiprotocol aspect of MPLS to switch MPLS frames between VPN endpoints while hiding the customer networks from the MPLS transport network that connects them. In other words, the intermediate transport network has no knowledge of the customer's IP networks, but it is still able to label-switch frames based on information it receives from MPLS Label Distribution Protocol (LDP) relationships.

You decide to model one of your current customer's connections and then show how MPLS VPNs can carry customer traffic through the provider network. The International Travel Agency currently uses your network to connect from its corporate headquarters to a remote branch office, so you choose this customer network to model in your demonstration.

First, set up the model of both the service provider's network and the agency's network. Then, use appropriate routing and forwarding techniques to set up a MPLS VPN between the provider edge routers to which the customer connects.

SP1, SP2, and SP3 represent a service-provider network, and HQ and BRANCH represent the International Travel Agency routers at their headquarters and at a branch site.

## Step 1: Configure Addressing

Configure the loopback interfaces with the addresses shown in Figure 4-2. Also, configure the serial interfaces shown in Figure 4-2. Set the clock rate on the appropriate interfaces and issue the **no shutdown** command on all physical interfaces. Verify that you have connectivity across the local subnet by using the **ping** command inside the service-provider domain. Wait to configure the interface on SP1 facing HQ and the interface on SP3 facing BRANCH. These are configured later:

```
SP1(config)# interface loopback 0
SP1(config-if)# ip address 10.0.1.1 255.255.255.255
SP1(config-if)# interface serial 0/0/0
SP1(config-if)# ip address 10.0.12.1 255.255.255.0
SP1(config-if)# no shutdown


---


SP2(config)# interface loopback 0
SP2(config-if)# ip address 10.0.2.1 255.255.255.255
SP2(config-if)# interface serial 0/0/0
SP2(config-if)# ip address 10.0.12.2 255.255.255.0
SP2(config-if)# no shutdown
SP2(config-if)# interface serial 0/0/1
SP2(config-if)# ip address 10.0.23.2 255.255.255.0
SP2(config-if)# clockrate 64000
SP2(config-if)# no shutdown


---


SP3(config)# interface loopback 0
SP3(config-if)# ip address 10.0.3.1 255.255.255.255
SP3(config-if)# interface serial 0/0/1
SP3(config-if)# ip address 10.0.23.3 255.255.255.0
SP3(config-if)# no shutdown
```

Configure customer sites HQ and BRANCH:

```
HQ(config)# interface loopback 0
HQ(config-if)# ip address 172.16.10.1 255.255.255.0
HQ(config-if)# interface fastethernet 0/0
HQ(config-if)# ip address 172.16.100.1 255.255.255.0
HQ(config-if)# no shutdown


---


BRANCH(config)# interface loopback 0
BRANCH(config-if)# ip address 172.16.20.1 255.255.255.0
BRANCH(config-if)# interface serial 0/0/0
BRANCH(config-if)# ip address 172.16.200.1 255.255.255.0
BRANCH(config-if)# clockrate 64000
BRANCH(config-if)# no shutdown
```

## Step 2: Configure Routing in the Service-Provider Domain

Your service-provider network uses OSPF as its routing protocol, advertising internal loopback interfaces and transit networks. Configure OSPF to model the service-provider domain. Add all the

interfaces addressed within the 10.0.0.0 major network into Area 0 of the OSPF process. You only need to configure OSPF in this manner on the service-provider routers—namely SP1, SP2, and SP3:

```
SP1(config)# router ospf 1
SP1(config-router)# network 10.0.0.0 0.255.255.255 area 0


---


SP2(config)# router ospf 1
SP2(config-router)# network 10.0.0.0 0.255.255.255 area 0


---


SP3(config)# router ospf 1
SP3(config-router)# network 10.0.0.0 0.255.255.255 area 0
```

Verify that all of your Open Shortest Path First (OSPF) adjacencies come up. OSPF adjacencies should form between SP1 and SP2 and between SP2 and SP3. If the adjacencies do not form, troubleshoot by checking your interface configuration, OSPF configuration, and physical connectivity.

What purpose does OSPF serve in the preceding configurations?

---

---

Consider that you will deploy BGP in the SP domain later in this lab using loopback addresses as the sources for BGP updates. Why do you need to deploy an Interior Gateway Protocol (IGP) in the SP domain?

---

---

---

---

---

---

### Step 3: Configure MPLS in the SP Domain

On all the service-provider routers, force MPLS to use the Loopback 0 interface as the router ID for LDP adjacencies. The loopback interface would be chosen by each router automatically, but it is advisable to force the ID so that the value is persistent through topology changes and reloads. To force LDP's selection of the loopback interface as the router ID, use the **mpls ldp router-id interface force** command in global configuration mode. Also, enable MPLS on all the physical interfaces in the MPLS domain with the **mpls ip** command:

```
SP1(config)# mpls ldp router-id loopback0 force
SP1(config)# interface serial0/0/0
SP1(config-if)# mpls ip


---


SP2(config)# mpls ldp router-id loopback0 force
SP2(config)# interface serial0/0/0
SP2(config-if)# mpls ip
SP2(config-if)# interface serial0/0/1
SP2(config-if)# mpls ip


---


SP3(config)# mpls ldp router-id loopback0 force
SP3(config)# interface serial0/0/1
SP3(config-if)# mpls ip
```

You should see console messages notifying you that the MPLS-enabled routers have become adjacent with each other via LDP. Verify that these adjacencies have formed using the **show mpls ldp neighbor** command:

SP1# **show mpls ldp neighbor**

```
Peer LDP Ident: 10.0.2.1:0; Local LDP Ident 10.0.1.1:0
TCP connection: 10.0.2.1.62676 - 10.0.1.1.646
State: Oper; Msgs sent/rcvd: 9/9; Downstream
Up time: 00:01:43
LDP discovery sources:
  Serial0/0/0, Src IP addr: 10.0.12.2
Addresses bound to peer LDP Ident:
  10.0.12.2      10.0.23.2      10.0.2.1
```

SP2# **show mpls ldp neighbor**

```
Peer LDP Ident: 10.0.1.1:0; Local LDP Ident 10.0.2.1:0
TCP connection: 10.0.1.1.646 - 10.0.2.1.62676
State: Oper; Msgs sent/rcvd: 10/10; Downstream
Up time: 00:02:03
LDP discovery sources:
  Serial0/0/0, Src IP addr: 10.0.12.1
Addresses bound to peer LDP Ident:
  10.0.12.1      10.0.1.1
```

```
Peer LDP Ident: 10.0.3.1:0; Local LDP Ident 10.0.2.1:0
TCP connection: 10.0.3.1.42919 - 10.0.2.1.646
State: Oper; Msgs sent/rcvd: 10/10; Downstream
Up time: 00:01:58
LDP discovery sources:
  Serial0/0/1, Src IP addr: 10.0.23.3
Addresses bound to peer LDP Ident:
  10.0.23.3      10.0.3.1
```

SP3# **show mpls ldp neighbor**

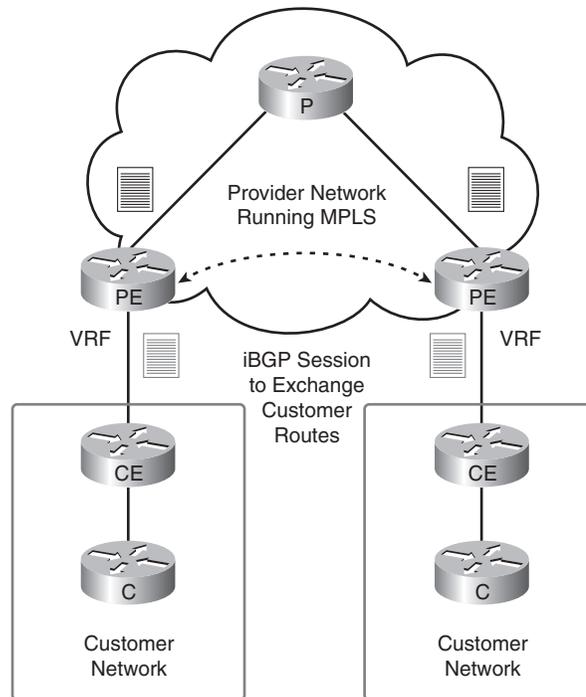
```
Peer LDP Ident: 10.0.2.1:0; Local LDP Ident 10.0.3.1:0
TCP connection: 10.0.2.1.646 - 10.0.3.1.42919
State: Oper; Msgs sent/rcvd: 10/10; Downstream
Up time: 00:02:08
LDP discovery sources:
  Serial0/0/1, Src IP addr: 10.0.23.2
Addresses bound to peer LDP Ident:
  10.0.12.2      10.0.23.2      10.0.2.1
```

## Step 4: Configure a VRF

An MPLS VPN is a Layer 3 VPN that allows packet routing through a MPLS core. This type of VPN provides a customer with connections to multiple sites through a service provider's network. The service provider not only provides the physical connection, but the ability to dynamically route between

the VPN endpoints. This is especially impressive when one considers that the customers might not be using globally unique Layer 3 addresses. For example, different customers can use private addresses, as defined by RFC 1918, but still use the same transit provider to route their specific endpoints without translation. As shown in Figure 4-3, the routers at the provider's edge run the same routing protocol as the customer's network and allow the customer offices to interface with the provider.

**Figure 4-3 MPLS VPNs Conceptual Diagram**



The standard model for MPLS VPNs uses the following designations:

- **Provider (P).** Routers owned by the service provider (SP) that act as label switch routers (LSR) to provide transit across the provider backbone. P routers do not carry customer routes in their routing tables.
- **Customer (C).** Routers owned by the customer that provide transit through the normal customer network.
- **Customer edge (CE).** The CE router is installed at the customer site. Depending on the business model of the Internet service provider (ISP), this router can be managed by the customer, the ISP, or both. The CE router connects to, and communicates with, the service-provider routers, and it allows the service provider to participate in customer routing.
- **Provider edge (PE).** Routers owned by the provider that actively participate in customer routing, guaranteeing optimum routing between customer sites. PE routers use a separate virtual routing table for each customer, which results in perfect isolation between customers.

Note that neither the C nor the CE routers need any special configuration. The P routers require only a simple MPLS LDP configuration.

In this lab, SP2 models the P router, and SP1 and SP3 model the PE routers. HQ and BRANCH are both CE routers with loopback networks to simulate connections to other C routers.

The PE routers control the entire MPLS VPN from end to end. You might be asking numerous relevant questions: How can a single router determine which routes in its table belong to the service provider's internal network, and which routes belong to each customer? How can the PE device allow customers to use existing networks, including private addressing without creating routing problems?

The answer to all of these questions lies in the ability of routers to maintain virtual routing and forwarding (VRF) instances. Each VRF uses and maintains its own routing information base (RIB) and Cisco Express Forwarding (CEF) table. Interfaces are either assigned to specific VRF instances or they use the default RIB and CEF tables. The VRF instance's RIB fulfills the role of control plane while the VRF's CEF table fulfills the role of the data forwarding plane. Routing protocols between the PE and CE routers populate the VRF RIB and CEF makes forwarding decisions based on the routes in the VRF RIB. When an IP packet arrives on an interface that has been associated with a VRF, the packet is routed according to the CEF table for that VRF instance. CEF is the only IP switching protocol supported for VRF, so CEF should be enabled globally with the **ip cef** command and on the interfaces associated with the VRF instance.

However, PE routers must now be connected through the provider network to perform routing and forwarding between customer sites. The most efficient and only scalable method to achieve this is to use the multiprotocol extensions to BGP (MP-BGP) that enable the provider network to carry routes for different routed protocols. PE routers establish iBGP sessions with other PEs in your carrier network to exchange for each VPN routes. This helps populate the VRF routing tables on each PE router with the VRF tables from other customer sites. CEF tables will be updated with the RIB information so that forwarding can occur between customer sites after the label-switched paths have been created through the provider network.

PE routers advertise routes that are part of their VPN by using a new traffic class to distinguish these routes from internal routes in the provider's network. BGP uses a new address family called VPNv4 to carry MPLS VPN routes to IPv4 networks. The VPNv4 address family is a 12-byte address consisting of an 8-byte route distinguisher (RD) and a 4-byte IPv4 address. The RD acts as a unique prefix when appended with the IPv4 address. Each VRF must have an RD for unique advertisement.

VRFs use the route target attribute to control the import and export of VPNv4 routes through iBGP. The route target is an extended BGP community that indicates which routes should be imported from MP-BGP into the VRF. Exporting a route target (RT) means that the exported VPNv4 route receives an additional BGP extended community—this is the route target—when the route is redistributed from the VRF RIB into MP-BGP. Importing an RT means that the received VPNv4 route from MP-BGP is checked for a matching extended community—this is the route target—with the ones in the configuration.

To configure a VRF instance on the PE routers, use the **ip vrf name** command in global configuration mode on SP1 and SP3. At the VRF configuration prompt, create a VRF named customer. Each VRF instance needs an RD and an RT. The RD and RT are each 8 bytes in length, with a colon separating 4 bytes on either side. There are various conventions for allocating RDs for MPLS VPNs; the most useful of which is ASN:nn. Another popular notation is IP address:nn. In each of these cases, nn represents an arbitrary value assigned by the network administrator. In this lab, use 100:1 as the RD. The RT is also an arbitrary 8-byte value used later in BGP.

Configure an RD of 100:1 and RT of 1:100 using the commands **rd ASN:nn** and **route-target {import | export | both} nn:nn**. In this case, you need to use the **both** keyword because you want PEs to import and export from that VRF:

```
SP1(config)# ip vrf customer
SP1(config-vrf)# rd 100:1
SP1(config-vrf)# route-target both 1:100
```

```
SP3(config)# ip vrf customer
SP3(config-vrf)# rd 100:1
SP3(config-vrf)# route-target both 1:100
```

Imagine that SP1 is running MP-BGP, and it receives a VPNv4 route with an RT of 100:100. Given the previous configuration, should BGP import the route into the *customer* VRF routing table?

---

After creating the VRFs, add interfaces to the VRF by using the interface-level **ip vrf forwarding name** command, where name is the VRF instance name. Use this command on the interfaces of SP1 and SP3 (the PE routers) facing the CE routers. Add the IP address, as shown in the figure, to those interfaces:

```
SP1(config)# interface fastethernet 0/0
SP1(config-if)# ip vrf forwarding customer
SP1(config-if)# ip address 172.16.100.254 255.255.255.0
SP1(config-if)# no shut


---


SP3(config)# interface serial 0/1/0
SP3(config-if)# ip vrf forwarding customer
SP3(config-if)# ip address 172.16.200.254 255.255.255.0
SP3(config-if)# no shutdown
```

You should now be able to ping across those the PE-CE links because you configured the other end of these links in Step 1. However, because these are not in the default routing table, you must use the **ping vrf name address** command. Because the VRF is transparent to the customer routers, you can use a traditional **ping** command when you ping from the C and CE routers:

```
SP1# ping vrf customer 172.16.100.1
```

Type escape sequence to abort.

```
Sending 5, 100-byte ICMP Echos to 172.16.100.1, timeout is 2 seconds:
```

```
!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

---

```
HQ# ping 172.16.100.254
```

Type escape sequence to abort.

```
Sending 5, 100-byte ICMP Echos to 172.16.100.254, timeout is 2 seconds:
```

```
!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/8 ms
```

---

```
SP3# ping vrf customer 172.16.200.1
```

Type escape sequence to abort.

```
Sending 5, 100-byte ICMP Echos to 172.16.200.1, timeout is 2 seconds:
```

```
!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/28/32 ms
```

---

```
BRANCH# ping 172.16.200.254
```

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 172.16.200.254, timeout is 2 seconds:

!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 28/28/28 ms

## Step 5: Configure EIGRP AS 1

The service provider by whom you are employed uses the BGP AS 100. Your customer, the International Travel Agency, uses the BGP AS 1. To keep the configuration logically consistent, use the AS number 100 for EIGRP and BGP in the provider's network, and use the AS number 1 for EIGRP and BGP in the customer's network. You configure EIGRP AS 1 on the PE routers from within the configuration of the global EIGRP AS 100.

On the customer routers, configure EIGRP AS 1 for the major network 172.16.0.0. Disable automatic summarization:

```
HQ(config)# router eigrp 1
HQ(config-router)# no auto-summary
HQ(config-router)# network 172.16.0.0
-----
BRANCH(config)# router eigrp 1
BRANCH(config-router)# no auto-summary
BRANCH(config-router)# network 172.16.0.0
```

Given only this information, will EIGRP immediately form any adjacencies?

---



---

On the PE routers, the configuration is more complex. Every IGP has a different method of configuring a VRF for it. To implement EIGRP for VRFs, start the EIGRP process by configuring EIGRP AS 100. Remember, this AS belongs to the provider and is not significant to the customer. If you were using EIGRP as the service provider's IGP instead of OSPF, you would configure your network statements at this point:

```
SP1(config)# router eigrp 100
-----
SP3(config)# router eigrp 100
```

Now, to configure EIGRP for an individual VRF instance, use the command **address-family ipv4 vrf name**, where *name* is the name of the VRF instance. Although each VPN must be logically separate from other IPv4 address spaces using VRF, this separation must extend not only to the routing table but also to the routing protocols. The **address-family** command creates a logical segment of a routing protocol and its routes and adjacencies to separate it from other sets of routes and adjacencies. In this case, you separate an EIGRP autonomous system from the EIGRP instance initiated with the **router eigrp 100** command. Networks learned via this new autonomous system are injected into the VRF routing table associated with the isolated EIGRP AS. It is important to note that these networks will not be advertised to any neighbors in EIGRP AS 100; it is completely separate from the rest of the EIGRP domain:

```
SP1(config-router)# address-family ipv4 vrf customer
SP1(config-router-af)# autonomous-system 1
SP1(config-router-af)# no auto-summary
SP1(config-router-af)# network 172.16.0.0
-----
```

```
SP3(config-router)# address-family ipv4 vrf customer
SP3(config-router-af)# autonomous-system 1
SP3(config-router-af)# no auto-summary
SP3(config-router-af)# network 172.16.0.0
```

On the PE routers, display the default routing table with the **show ip route** command. Notice that the PE routers do not possess any routes from the 172.16.0.0/16 major network in the default routing table. Display the VRF routing table with the **show ip route vrf name** command, where *name* is the VRF instance name:

```
SP1# show ip route
```

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route
```

Gateway of last resort is not set

```
10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
C       10.0.12.0/24 is directly connected, Serial0/0/0
O       10.0.3.1/32 [110/129] via 10.0.12.2, 05:29:59, Serial0/0/0
O       10.0.2.1/32 [110/65] via 10.0.12.2, 05:29:59, Serial0/0/0
C       10.0.1.1/32 is directly connected, Loopback0
O       10.0.23.0/24 [110/128] via 10.0.12.2, 05:29:59, Serial0/0/0
```

```
SP1# show ip route vrf customer
```

Routing Table: customer

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route
```

Gateway of last resort is not set

```
172.16.0.0/24 is subnetted, 2 subnets
D       172.16.10.0 [90/156160] via 172.16.100.1, 00:03:29, FastEthernet0/0
C       172.16.100.0 is directly connected, FastEthernet0/0
```

---

```
SP3# show ip route
```

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
```

```

D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2
i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
ia - IS-IS inter area, * - candidate default, U - per-user static route
o - ODR, P - periodic downloaded static route

```

Gateway of last resort is not set

```

10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
O    10.0.12.0/24 [110/128] via 10.0.23.2, 05:30:42, Serial0/0/1
C    10.0.3.1/32 is directly connected, Loopback0
O    10.0.2.1/32 [110/65] via 10.0.23.2, 05:30:42, Serial0/0/1
O    10.0.1.1/32 [110/129] via 10.0.23.2, 05:30:42, Serial0/0/1
C    10.0.23.0/24 is directly connected, Serial0/0/1

```

---

```

SP3# show ip route vrf customer

```

Routing Table: customer

```

Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2
i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
ia - IS-IS inter area, * - candidate default, U - per-user static route
o - ODR, P - periodic downloaded static route

```

Gateway of last resort is not set

```

172.16.0.0/24 is subnetted, 2 subnets
C    172.16.200.0 is directly connected, Serial0/1/0
D    172.16.20.0 [90/2297856] via 172.16.200.1, 00:02:06, Serial0/1/0

```

The SP1 and HQ routers do not possess routes to the customer networks on SP3 and BRANCH and vice versa. Explain why this occurs, even though EIGRP adjacencies have formed.

---



---



---

## Step 6: Configure BGP

Now that the PE routers are routing to the CE routers over VRF tables, you can set up the PE routers to exchange routes through BGP. First, configure BGP between SP1 and SP3 and have them peer between their loopback addresses. Synchronization should be disabled by default on newer IOS releases. If synchronization is not already disabled, explicitly disable it by using the **no synchronization** command. For more information on configuring BGP, refer to CCNP1:

```
SP1(config)# router bgp 100
SP1(config-router)# neighbor 10.0.3.1 remote-as 100
SP1(config-router)# neighbor 10.0.3.1 update-source loopback0


---


SP3(config)# router bgp 100
SP3(config-router)# neighbor 10.0.1.1 remote-as 100
SP3(config-router)# neighbor 10.0.1.1 update-source loopback0
```

To configure the exchange of VPNv4 routes over BGP, use the **address-family vpnv4** command. At the address family prompt, activate the BGP neighbor for this address family with **neighbor address activate** command. Activating a neighbor for an address family allows BGP to send routes to and receive routes from the designated neighbor using the specified address family. By default, neighbors are only activated for IPv4.

The RTs are translated as extended BGP communities, so you must allow SP1 and SP3 to send both standard and extended communities over MP-BGP using the **neighbor address send-community both** command. The adjacencies might flap (temporarily go down and then come back up) when you activate the address family.

```
SP1(config-router)# address-family vpnv4
SP1(config-router-af)# neighbor 10.0.3.1 activate
SP1(config-router-af)# neighbor 10.0.3.1 send-community both
SP1(config-router-af)# exit


---


SP3(config-router)# address-family vpnv4
SP3(config-router-af)# neighbor 10.0.1.1 activate
SP3(config-router-af)# neighbor 10.0.1.1 send-community both
SP3(config-router-af)# exit
```

Finally, you need to configure BGP to redistribute the EIGRP routes in the VRF RIB into the BGP protocol so that these routes are advertised to the remote PE. Under the main BGP configuration prompt, enter another address family associated only with the routing table for the VRF *customer*. Redistribute the EIGRP routes that are associated with this VRF into BGP:

```
SP1(config-router)# address-family ipv4 vrf customer
SP1(config-router-af)# redistribute eigrp 1
SP1(config-router-af)# exit
SP1(config-router)# exit


---


SP3(config-router)# address-family ipv4 vrf customer
SP3(config-router-af)# redistribute eigrp 1
SP3(config-router-af)# exit
SP3(config-router)# exit
```

Based on the previous configuration, will SP1's VRF RIB contain the 172.16.20.0/24 route that was originated by EIGRP on BRANCH? Explain.

---

---

Will HQ learn the same routes via EIGRP? Explain.

---

---

Do you expect to see the redistributed routes as internal or external EIGRP routes on the CE routers? Explain.

---



---



---



---



---

Enter the EIGRP instance that contains the VRF configuration on SP1 and SP3, and configure it to redistribute BGP routes. Because you are redistributing into EIGRP from BGP, the metrics are not comparable. Add a seed metric with a bandwidth of 64 kbps, 100 microseconds, reliability of 255/255, load of 1/255, and MTU of 1500 bytes:

```
SP1(config)# router eigrp 100
SP1(config-router)# address-family ipv4 vrf customer
SP1(config-router-af)# redistribute bgp 100 metric 64 1000 255 1 1500
SP3(config)# router eigrp 100
SP3(config-router)# address-family ipv4 vrf customer
SP3(config-router-af)# redistribute bgp 100 metric 64 1000 255 1 1500
```

## Step 7: Investigate Control Plane Operation

Remember that MPLS differentiates the control plane from the forwarding plane. The control plane, represented by the routing table (the RIB) and the routing protocols, must operate so that the VRF routes reach remote PEs and are installed as necessary in the VRF routing tables. Not only the prefixes, but also the accompanying metrics and tags are important to the reconstruction of the route at the remote PE. Fortunately, MP-BGP allows you to send these metrics in the Network Layer Reachability Information (NLRI).

Through this step and Step 8, you investigate the routing and forwarding information associated with the route to 172.16.20.0/24.

Verify that the routes have propagated to the remote PE routers. Issue the **show ip route vrf name** command to see the VRF RIB. Notice the source of the routes on the PE routers:

```
SP1# show ip route vrf customer
```

```
Routing Table: customer
```

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
```

```
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
```

```
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
```

```
       E1 - OSPF external type 1, E2 - OSPF external type 2
```

```
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
```

```
       ia - IS-IS inter area, * - candidate default, U - per-user static route
```

```
       o - ODR, P - periodic downloaded static route
```

```
Gateway of last resort is not set
```

172.16.0.0/24 is subnetted, 4 subnets

```
B    172.16.200.0 [200/0] via 10.0.3.1, 00:06:44
B    172.16.20.0 [200/2297856] via 10.0.3.1, 00:06:44
D    172.16.10.0 [90/156160] via 172.16.100.1, 00:17:34, FastEthernet0/0
C    172.16.100.0 is directly connected, FastEthernet0/0
```

SP3# **show ip route vrf customer**

Routing Table: customer

Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP

D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area

N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2

E1 - OSPF external type 1, E2 - OSPF external type 2

i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2

ia - IS-IS inter area, \* - candidate default, U - per-user static route

o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

172.16.0.0/24 is subnetted, 4 subnets

```
C    172.16.200.0 is directly connected, Serial0/1/0
D    172.16.20.0 [90/2297856] via 172.16.200.1, 16:47:37, Serial0/1/0
B    172.16.10.0 [200/156160] via 10.0.1.1, 00:17:28
B    172.16.100.0 [200/0] via 10.0.1.1, 00:17:28
```

You might ask, “Why does the source of the route to 172.16.20.0/24 on SP1 point to 10.0.3.1, because that address would be routed based on the default routing table?” Consider that, when an internally generated route is sent to an iBGP peer, BGP sets the next-hop attribute to be the advertising router. In this case, SP3 generates the route in BGP by redistribution. The BGP peers are communicating between loopback interfaces, so the next-hop is set to the IP address of the BGP peer’s source interface. Thus, the VRF RIB points to an interface that must be reached through the default global RIB. We will investigate the forwarding for packets destined for these networks in the next step.

On the CE routers, issue the **show ip route** command to see a full routing table:

HQ# **show ip route**

Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP

D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area

N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2

E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP

i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2

ia - IS-IS inter area, \* - candidate default, U - per-user static route

o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

```

172.16.0.0/24 is subnetted, 4 subnets
D       172.16.200.0
        [90/2172416] via 172.16.100.254, 00:05:17, FastEthernet0/0
D       172.16.20.0 [90/2300416] via 172.16.100.254, 00:05:17, FastEthernet0/0
C       172.16.10.0 is directly connected, Loopback0
C       172.16.100.0 is directly connected, FastEthernet0/0

```

BRANCH# **show ip route**

```

Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
        D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
        N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
        E1 - OSPF external type 1, E2 - OSPF external type 2
        i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
        ia - IS-IS inter area, * - candidate default, U - per-user static route
        o - ODR, P - periodic downloaded static route

```

Gateway of last resort is not set

```

172.16.0.0/24 is subnetted, 4 subnets
C       172.16.200.0 is directly connected, Serial0/0/0
C       172.16.20.0 is directly connected, Loopback0
D       172.16.10.0 [90/2300416] via 172.16.200.254, 00:02:02, Serial0/0/0
D       172.16.100.0 [90/2172416] via 172.16.200.254, 00:02:02, Serial0/0/0

```

On both the CE and PE routers, notice that the routes you redistributed from BGP into EIGRP are *internal* EIGRP routes because BGP preserves features of the EIGRP route while advertising the route to the other PEs. The PE encodes as much EIGRP information as possible into new extended communities—TLV tuples (type, length, value)—to preserve route characteristics through the VPN. This enables the remote PE router to reconstruct the EIGRP route with all of its characteristics, including the metric components, AS, TAG, and, for external routes, the remote AS number, the remote ID, the remote protocol, and the remote metric. These are the EIGRP characteristics of a prefix that you can find in the topology table. If the EIGRP-advertised route is internal, the route is advertised as an internal route into the remote site if the destination AS matches the source AS carried by the BGP extended community attributes.

Display information on the VPNv4 BGP routes on SP1 with the **show bgp vpnv4 unicast all** command:

```
SP1# show bgp vpnv4 unicast all
```

```
BGP table version is 9, local router ID is 10.0.1.1
```

```
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale
```

```
Origin codes: i - IGP, e - EGP, ? - incomplete
```

Network	Next Hop	Metric	LocPrf	Weight	Path
Route Distinguisher: 100:1 (default for vrf customer)					
*> 172.16.10.0/24	172.16.100.1	156160		32768	?
*>i172.16.20.0/24	10.0.3.1	2297856	100	0	?

```
*> 172.16.100.0/24 0.0.0.0 0 32768 ?
*>i172.16.200.0/24 10.0.3.1 0 100 0 ?
```

---

SP3# **show bgp vpnv4 unicast all**

BGP table version is 9, local router ID is 10.0.3.1

Status codes: s suppressed, d damped, h history, \* valid, > best, i - internal,  
r RIB-failure, S Stale

Origin codes: i - IGP, e - EGP, ? - incomplete

Network	Next Hop	Metric	LocPrf	Weight	Path
Route Distinguisher: 100:1 (default for vrf customer)					
*>i172.16.10.0/24	10.0.1.1	156160	100	0	?
*> 172.16.20.0/24	172.16.200.1	2297856		32768	?
*>i172.16.100.0/24	10.0.1.1	0	100	0	?
*> 172.16.200.0/24	0.0.0.0	0		32768	?

Notice that the metric (MED value) in BGP is the metric advertised through EIGRP for that route as well.

What does the value of the NEXT-HOP attribute for the 172.16.200.0/24 network on SP3 indicate?

---

What is the value of the BGP NEXT-HOP attribute for the 172.16.20.0/24 route on SP1?

---

By which routing protocol, and from which router, was the route to 10.0.3.1/32 installed in the default routing table on SP1?

---

View more specific detail on a particular prefix by using **show bgp vpnv4 unicast all ip-address** command. Notice that the MPLS label information is included. Execute this on both of the PEs. Remember that SP3 is advertising the 172.16.20.0/24 prefix through BGP, while SP1 is receiving the route through BGP NLRI:

SP1# **show bgp vpnv4 unicast all 172.16.20.0/24**

BGP routing table entry for 100:1:172.16.20.0/24, version 15

Paths: (1 available, best #1, table customer)

Flag: 0x820

Not advertised to any peer

Local

10.0.3.1 (metric 129) from 10.0.3.1 (10.0.3.1)

Origin incomplete, metric 2297856, localpref 100, valid, internal, best

Extended Community: RT:1:100

Cost:pre-bestpath:128:2297856 (default-2145185791) 0x8800:32768:0

0x8801:1:640000 0x8802:65281:1657856 0x8803:65281:1500

mpls labels in/out no-label/20

---

SP3# **show bgp vpnv4 unicast all 172.16.20.1**

BGP routing table entry for 100:1:172.16.20.0/24, version 15

Paths: (1 available, best #1, table customer)

```

Advertised to update-groups:
  1
Local
172.16.200.1 from 0.0.0.0 (10.0.3.1)
  Origin incomplete, metric 2297856, localpref 100, weight 32768, valid, sourced,
  best
  Extended Community: RT:1:100
    Cost:pre-bestpath:128:2297856 (default-2145185791) 0x8800:32768:0
    0x8801:1:640000 0x8802:65281:1657856 0x8803:65281:1500
  mpls labels in/out 20/nolabel

```

Notice that multiple values are in the BGP extended communities. Recall that BGP sends the route information in NLRI as extended communities. These values are TLVs, indicating such EIGRP attributes as the TAG, AS number, bandwidth, delay, reliability, load, MTU, and hop count.

Why is the origin code “incomplete?”

---



---

What type of attribute carries the route target information in MP-BGP NLRI?

---



---

Notice the MPLS labels indicated for this BGP route. The in-label of nolabel on SP1 indicates that SP1 is not advertising a label for the prefix 172.16.20.0/24. The out-label of 21 is advertised by SP3 and received by SP1. This label is significant only on the path between SP1 and SP3. This label has been allocated by BGP on SP3.

View the list of MPLS labels that are being used with BGP using **show bgp vpnv4 unicast all labels**:

SP1# **show bgp vpnv4 unicast all labels**

Network	Next Hop	In label/Out label
Route Distinguisher: 100:1 (customer)		
172.16.10.0/24	172.16.100.1	19/nolabel
172.16.20.0/24	10.0.3.1	nolabel/20
172.16.100.0/24	0.0.0.0	20/aggregate(customer)
172.16.200.0/24	10.0.3.1	nolabel/19

SP3# **show bgp vpnv4 unicast all labels**

Network	Next Hop	In label/Out label
Route Distinguisher: 100:1 (customer)		
172.16.10.0/24	10.0.1.1	nolabel/19
172.16.20.0/24	172.16.200.1	20/nolabel
172.16.100.0/24	10.0.1.1	nolabel/20
172.16.200.0/24	0.0.0.0	19/aggregate(customer)

How has SP1 learned the VPN label, label 20?

---



---

Will SP1 or SP2 learn the label via LDP?

---



---

Has the P router SP2 learned about label 20 from SP3? Explain.

---

Finally, display the route attributes for the same prefix, 172.16.20.0/24, in the EIGRP topology table on SP1 with the **show ip eigrp vrf customer topology ip-prefix/mask** command. Verify this against the originator of the EIGRP route in BGP, SP3:

```
SP1# show ip eigrp vrf customer topology 172.16.20.0/24
```

```
IP-EIGRP (AS 1): Topology entry for 172.16.20.0/24
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 2297856
  Routing Descriptor Blocks:
    10.0.3.1, from VPNv4 Sourced, Send flag is 0x0
      Composite metric is (2297856/0), Route is Internal (VPNv4 Sourced)
      Vector metric:
        Minimum bandwidth is 1544 Kbit
        Total delay is 25000 microseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 1
```

---

```
SP3# show ip eigrp vrf customer topology 172.16.20.0/24
```

```
IP-EIGRP (AS 1): Topology entry for 172.16.20.0/24
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 2297856
  Routing Descriptor Blocks:
    172.16.200.1 (Serial0/1/0), from 172.16.200.1, Send flag is 0x0
      Composite metric is (2297856/128256), Route is Internal
      Vector metric:
        Minimum bandwidth is 1544 Kbit
        Total delay is 25000 microseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 1
```

There is absolutely no difference in the EIGRP route parameters between SP1 and SP3. BGP encodes and decodes the information on the PE routers with no changes.

Remember that SP2—a P router—has no knowledge of individual routes in the VRF tables on the PE routers. You can verify this with the **show** commands previously performed:

```
SP2# show ip route
```

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
```

```

ia - IS-IS inter area, * - candidate default, U - per-user static route
o - ODR, P - periodic downloaded static route

```

```

Gateway of last resort is not set

```

```

10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
C    10.0.12.0/24 is directly connected, Serial0/0/0
O    10.0.3.1/32 [110/65] via 10.0.23.3, 1d00h, Serial0/0/1
C    10.0.2.1/32 is directly connected, Loopback0
O    10.0.1.1/32 [110/65] via 10.0.12.1, 1d00h, Serial0/0/0
C    10.0.23.0/24 is directly connected, Serial0/0/1

```

```

SP2# show ip route vrf customer

```

```

% IP routing table customer does not exist

```

Ping between the CE routers to verify connectivity through the MPLS VPN:

```

HQ# ping 172.16.20.1

```

```

Type escape sequence to abort.

```

```

Sending 5, 100-byte ICMP Echos to 172.16.20.1, timeout is 2 seconds:

```

```

!!!!

```

```

Success rate is 100 percent (5/5), round-trip min/avg/max = 84/89/93 ms

```

```

BRANCH# ping 172.16.10.1

```

```

Type escape sequence to abort.

```

```

Sending 5, 100-byte ICMP Echos to 172.16.10.1, timeout is 2 seconds:

```

```

!!!!

```

```

Success rate is 100 percent (5/5), round-trip min/avg/max = 84/86/88 ms

```

## Step 8: Investigate Forwarding Plane Operation

Recall that MPLS has two tables: the Label Information Base (LIB) and the Label Forwarding Information Base (LFIB). Normally, LDP-allocated labels are advertised to LDP peers. BGP-allocated labels are advertised to BGP peers. BGP-allocated labels will be used by BGP peers as an MPLS label on packets destined for that network through the VPN. The BGP-allocated labels are only significant to the ingress and egress routers. P routers that are not BGP peers with the PE routers will not see the VPN label for the networks known by BGP.

To traverse the MPLS cloud, the packets need to be label-switched at every hop based on advertised labels. To ensure that VPN packets that reach the egress PE have the MPLS label needed to switch the packets after they arrive, the labels are stacked at the ingress PE. However, the packet still needs to be sent along the label-switched path.

Recall that the VRF RIB's next hop for the networks is known via the VPN to the loopback on the egress PE. CEF uses the inuse label for the BGP next hop as the outermost label for packets traveling through the MPLS VPN. First, however, CEF must push on the VPN label that will be used at the egress PE. Thus, CEF stacks the label in a sequential manner so that the VPN label is available at the

egress PE, but the label to traverse the label-switched path through the P routers is pushed as the outermost label.

Take some time to study and understand the details of how this is possible. BGP, LDP, CEF, the LFIB, and the provider's IGP are all involved in the use of MPLS labels as a VPN technology.

Once BGP learns the MPLS label to use as the VPN label, this information is entered into the CEF forwarding table on the ingress PE. Display the CEF forwarding entry for 172.16.20.0/24 on SP1 with the **show ip cef vrf name ip-address** command:

```
SP1# show ip cef vrf customer 172.16.20.0
172.16.20.0/24, version 12, epoch 0, cached adjacency to Serial0/0/0
0 packets, 0 bytes
tag information set
  local tag: VPN-route-head
  fast tag rewrite with Se0/0/0, point2point, tags imposed: {16 20}
via 10.0.3.1, 0 dependencies, recursive
  next hop 10.0.12.2, Serial0/0/0 via 10.0.3.1/32
  valid cached adjacency
  tag rewrite with Se0/0/0, point2point, tags imposed: {16 20}
```

CEF resolves the recursive lookup to the BGP next hop. Based on the labels learned by LDP, CEF might or might not apply the forwarding label to reach 10.0.3.1/32. In this case, LDP on SP2 has advertised a forwarding label to SP1. View the labels advertised to SP1 via LDP using the **show mpls ip binding** command:

```
SP1# show mpls ip binding
10.0.1.1/32
  in label:    imp-null
  out label:   17      lsr: 10.0.2.1:0
10.0.2.1/32
  in label:    16
  out label:   imp-null lsr: 10.0.2.1:0      inuse
10.0.3.1/32
  in label:    17
  out label:   16      lsr: 10.0.2.1:0      inuse
10.0.12.0/24
  in label:    imp-null
  out label:   imp-null lsr: 10.0.2.1:0
10.0.23.0/24
  in label:    18
  out label:   imp-null lsr: 10.0.2.1:0      inuse
```

CEF pushes the label of 20 onto the packet first, and then pushes the outer label of 16. The CEF forwarding table decides which path to use based, of course, on the default RIB. The route has been installed in the RIB by OSPF. Thus, the ingress PE imposes two labels in the sequence {16, 20}, as shown in this CEF forwarding table.

Because the incoming VPN packets from SP1 are encapsulated in MPLS frames, SP2 acts according to the directives in its LFIB. SP2 is also the penultimate hop in the label-switched path from SP1 to SP3's loopback interface, and therefore pops the outermost label from the MPLS frame. Display the LFIB with the **show mpls forwarding-table** command:

```
SP2# show mpls forwarding-table
```

Local tag	Outgoing tag or VC	Prefix or Tunnel Id	Bytes tag switched	Outgoing interface	Next Hop
16	Pop tag	10.0.3.1/32	5175	Se0/0/1	point2point
17	Pop tag	10.0.1.1/32	8079	Se0/0/0	point2point

Notice that the LFIB does not care whether there is an inner label or not; it simply performs the operation specified in the column labeled Outgoing, tag or VC.

If you enable MPLS packet debugging on SP2 using **debug mpls packets**, and then issue a ping from one CE to the other, you can see the MPLS packets being label-switched. The ICMP packets are forwarded inside MPLS frames through SP2. Notice in the debug output that each ICMP echo request receives a reply, which is label-switched on its return path through the MPLS network. When you are done, disable debugging:

```
SP2# debug mpls packets
```

```
MPLS packet debugging is on
```

```
HQ# ping 172.16.20.1
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 172.16.20.1, timeout is 2 seconds:
```

```
!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 84/87/92 ms
```

```
SP2#
```

```
*Feb 3 20:55:57.422: MPLS: Se0/0/0: recvd: CoS=0, TTL=254, Label(s)=16/20
```

```
*Feb 3 20:55:57.422: MPLS: Se0/0/1: xmit: CoS=0, TTL=253, Label(s)=20
```

```
!
```

```
! These 2 messages indicate the label-switching of the ICMP echo request
```

```
!
```

```
*Feb 3 20:55:57.478: MPLS: Se0/0/1: recvd: CoS=0, TTL=254, Label(s)=17/20
```

```
*Feb 3 20:55:57.478: MPLS: Se0/0/0: xmit: CoS=0, TTL=253, Label(s)=20
```

```
!
```

```
! These 2 messages indicate the label-switching of the ICMP echo reply
```

```
!
```

```
*Feb 3 20:55:57.510: MPLS: Se0/0/0: recvd: CoS=0, TTL=254, Label(s)=16/20
```

```
*Feb 3 20:55:57.510: MPLS: Se0/0/1: xmit: CoS=0, TTL=253, Label(s)=20
```

```
*Feb 3 20:55:57.566: MPLS: Se0/0/1: recvd: CoS=0, TTL=254, Label(s)=17/20
```

```
*Feb 3 20:55:57.566: MPLS: Se0/0/0: xmit: CoS=0, TTL=253, Label(s)=20
```

```
*Feb 3 20:55:57.598: MPLS: Se0/0/0: recvd: CoS=0, TTL=254, Label(s)=16/20
```

```
*Feb 3 20:55:57.598: MPLS: Se0/0/1: xmit: CoS=0, TTL=253, Label(s)=20
```

```
*Feb 3 20:55:57.654: MPLS: Se0/0/1: recvd: CoS=0, TTL=254, Label(s)=17/20
```

```
*Feb 3 20:55:57.654: MPLS: Se0/0/0: xmit: CoS=0, TTL=253, Label(s)=20
```

```
*Feb 3 20:55:57.686: MPLS: Se0/0/0: recvd: CoS=0, TTL=254, Label(s)=16/20
*Feb 3 20:55:57.686: MPLS: Se0/0/1: xmit: CoS=0, TTL=253, Label(s)=20
*Feb 3 20:55:57.742: MPLS: Se0/0/1: recvd: CoS=0, TTL=254, Label(s)=17/20
*Feb 3 20:55:57.742: MPLS: Se0/0/0: xmit: CoS=0, TTL=253, Label(s)=20
*Feb 3 20:55:57.774: MPLS: Se0/0/0: recvd: CoS=0, TTL=254, Label(s)=16/20
*Feb 3 20:55:57.774: MPLS: Se0/0/1: xmit: CoS=0, TTL=253, Label(s)=20
*Feb 3 20:55:57.830: MPLS: Se0/0/1: recvd: CoS=0, TTL=254, Label(s)=17/20
*Feb 3 20:55:57.830: MPLS: Se0/0/0: xmit: CoS=0, TTL=253, Label(s)=20
```

```
SP2# undebug all
```

All possible debugging has been turned off

Continue tracing the label-switched path through the provider network to the egress PE, SP3.

Based on which forwarding table will the VPN packet be switched at SP3? Explain.

Display the MPLS LFIB on SP3 using the **show mpls forwarding-table** command that you used on SP2 previously:

```
SP3# show mpls forwarding-table
```

Local tag	Outgoing tag or VC	Prefix or Tunnel Id	Bytes tag switched	Outgoing interface	Next Hop
16	Pop tag	10.0.12.0/24	0	Se0/0/1	point2point
17	Pop tag	10.0.2.1/32	0	Se0/0/1	point2point
18	17	10.0.1.1/32	0	Se0/0/1	point2point
19	Aggregate	172.16.200.0/24[V] \	2704		
20	Untagged	172.16.20.0/24[V] 2704	2704	Se0/1/0	point2point

Notice that SP3 forwards the decapsulated IP packet untagged to the Serial 0/1/0 egress interface because it was received with a label of 20. This is the label that BGP advertised to SP1. SP1's CEF forwarding table encapsulated the IP packets within two MPLS labels {16 20} and then forwarded the packet to SP2.

## Conclusion

Issue the **traceroute** command from one CE to another to find that it is going through multiple Layer 3 hops. This is an important debugging tool because it can also be issued from a PE router with reference to a VRF:

```
HQ# traceroute 172.16.20.1
```

Type escape sequence to abort.

Tracing the route to 172.16.20.1

```

1 172.16.100.254 0 msec 0 msec 0 msec
2 10.0.12.2 126 msec 117 msec 126 msec
3 172.16.200.254 59 msec 50 msec 50 msec
4 172.16.200.1 50 msec 42 msec *
```

Fill in Table 4-1, tracing the path of packets from 172.16.100.1 to 172.16.20.1 to trace the packet's path.

**Table 4-1 Path of Packets Through the Network**

Router	Incoming (MPLS/IP)	Outgoing (MPLS/IP)	Switched By (CEF/LFIB)	Incoming Label(s)	Outgoing Label(s)
HQ	—			—	—
SP1				—	
SP2					
SP3					—
BRANCH		—		—	—

Given the following output on each of the routers, trace the return path from 172.16.20.1 to 172.16.100.1 by filling in the chart shown in Table 4-2:

BRANCH# **show ip cef 172.16.100.1**

```

172.16.100.0/24, version 22, epoch 0, cached adjacency to Serial0/0/0
0 packets, 0 bytes
  via 172.16.200.254, Serial0/0/0, 0 dependencies
    next hop 172.16.200.254, Serial0/0/0
    valid cached adjacency
```

SP3# **show ip cef vrf customer 172.16.100.1**

```

172.16.100.0/24, version 6, epoch 0, cached adjacency to Serial0/0/1
0 packets, 0 bytes
  tag information set
    local tag: VPN-route-head
    fast tag rewrite with Se0/0/1, point2point, tags imposed: {17 20}
  via 10.0.1.1, 0 dependencies, recursive
    next hop 10.0.23.2, Serial0/0/1 via 10.0.1.1/32
    valid cached adjacency
    tag rewrite with Se0/0/1, point2point, tags imposed: {17 20}
```

SP2# **show mpls forwarding-table**

Local tag	Outgoing tag or VC	Prefix or Tunnel Id	Bytes tag switched	Outgoing interface	Next Hop
16	Pop tag	10.0.3.1/32	15601	Se0/0/1	point2point
17	Pop tag	10.0.1.1/32	25413	Se0/0/0	point2point

SP1# **show mpls forwarding-table**

Local tag	Outgoing tag or VC	Prefix or Tunnel Id	Bytes tag switched	Outgoing interface	Next Hop
16	Pop tag	10.0.2.1/32	0	Se0/0/0	point2point

17	16	10.0.3.1/32	0	Se0/0/0	point2point
18	Pop tag	10.0.23.0/24	0	Se0/0/0	point2point
19	Untagged	172.16.10.0/24[V]	0	Fa0/0	172.16.100.1
20	Aggregate	172.16.100.0/24[V]	\0		

**Table 4-2 Return Path of Packets Through the Network**

Router	Incoming (MPLS/IP)	Outgoing (MPLS/IP)	Switched By (CEF/LFIB)	Incoming Label(s)	Outgoing Label(s)
BRANCH	—			—	—
SP3				—	
SP2					
SP1					—
HQ		—		—	—