MPLS

Implementing Cisco MPLS

Version 2.2

Lab Guide

Text Part Number: 97-2391-01

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.comWebsiteatwww.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

Lab Guide

Overview

This guide presents the instructions and other information concerning the activities for this course. You can find the solutions in the activity Answer Key.

Outline

This guide includes these activities:

- Lab 2-1: Establishing the Service Provider IGP Routing Environment
- Lab 3-1: Establishing the Core MPLS Environment
- Lab 5-1: Configuring Initial MPLS VPN Setup
- Lab 5-2: Running EIGRP Between PE and CE Routers
- Lab 5-3: Running OSPF Between PE and CE Routers
- Lab 5-4: Running BGP Between PE and CE Routers
- Lab 6-1: Establishing Overlapping VPNs
- Lab 6-2: Merging Service Providers
- Lab 6-3: Establishing a Common Services VPN
- Lab 7-1: Establishing Central Site Internet Connectivity with an MPLS VPN
- Lab 8-1: Implementing Basic MPLS TE
- Answer Key

Lab 2-1: Establishing the Service Provider IGP Routing Environment

Complete this lab activity to practice what you learned in the related module.

Activity Objective

In this activity, you will use the tasks and commands necessary to implement the service provider IGP and routing environment. After completing this activity, you will be able to meet these objectives:

- Verify the service provider IP addressing scheme, DLCI assignment, and interface status
- Enable the service provider IGP and configure appropriate IP addressing

Visual Objective

The figure illustrates what you will accomplish in this activity. This activity contains information about your laboratory setup, details of the physical and logical connectivity in the laboratory, and information about the addressing scheme and IGP routing. The class will be divided into service providers, or SPs, (where *x* represents your assigned SP number). Each SP will contain the router types as defined in the table.

The names of all routers in your SP follow the naming convention detailed in this table.

Router Naming Convention

Router Role	Description
P (provider)	P*x*1 and P*x*2 are core routers in the network of the provider.
PE(provider edge)	PE*x*1 and PE*x*2 are provider edge routers connecting from the provider to the customer network.
CE(customer edge)	CE*x*1A and CE*x*2A and CE*x*1B and CE*x*2B are customer edge routers for customer A and customer B, respectively.

SP numbering will be provided by your instructor. SP 3, 4, 5, and 6 will be used for this course.

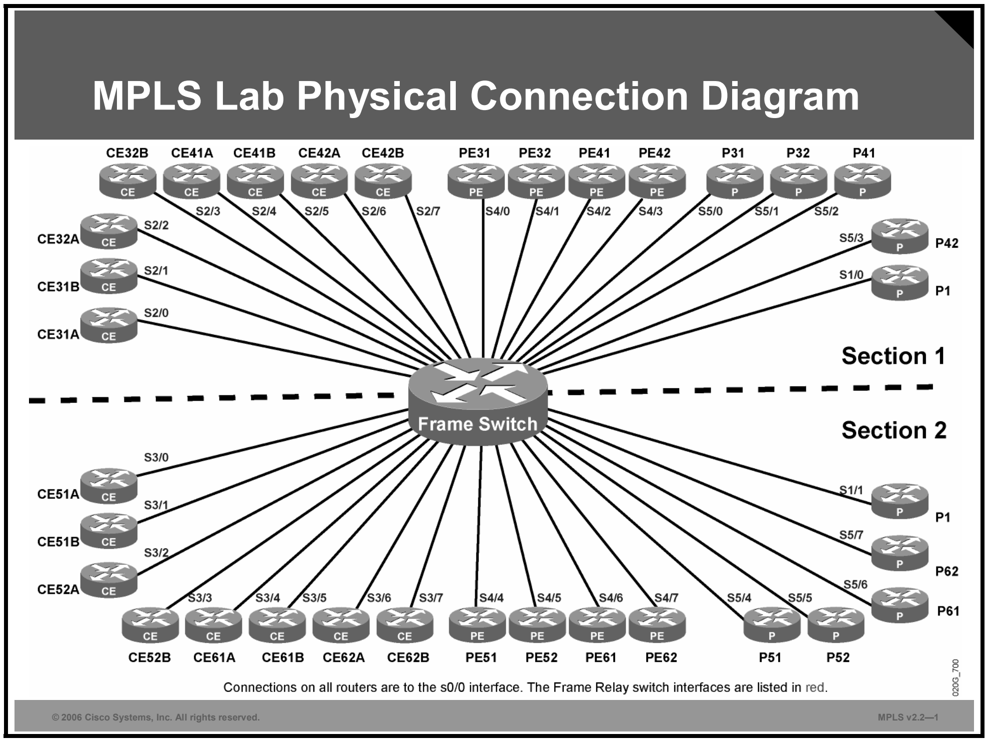

Physical connectivity has been provided by preconfigured PVCs defined by their respective DLCIs. The first serial interface of each router (P, PE, and CE) is connected to a Frame Relay switch. The DLCI values for all Frame Relay virtual circuits are shown in the DLCI identification table and the logical connection diagram figure. The subinterface number matches the DLCI values for all Frame Relay virtual circuits.

DLCI Identification

Source Router	Destination Router	DLCI	Interface
CE*x*1A	PE*x*1	101	S0/0.101
CE*x*1B	PE*x*1	102	S0/0.102
CE*x*2A	PE*x*2	101	S0/0.101
CE*x*2B	PE*x*2	102	S0/0.102
PE*x*1	CE*x*1A	101	S0/0.101
PE*x*1	CE*x*1B	102	S0/0.102
PE*x*1	P*x*1	111	S0/0.111
PE*x*2	CE*x*2A	101	S0/0.101
PE*x*2	CE*x*2B	102	S0/0.102
PE*x*2	P*x*2	111	S0/0.111
P*x*1	PE*x*1	111	S0/0.111
P*x*1	P*x*2	112	S0/0.112
P*x*2	PE*x*2	111	S0/0.111
P*x*2	P*x*1	112	S0/0.112

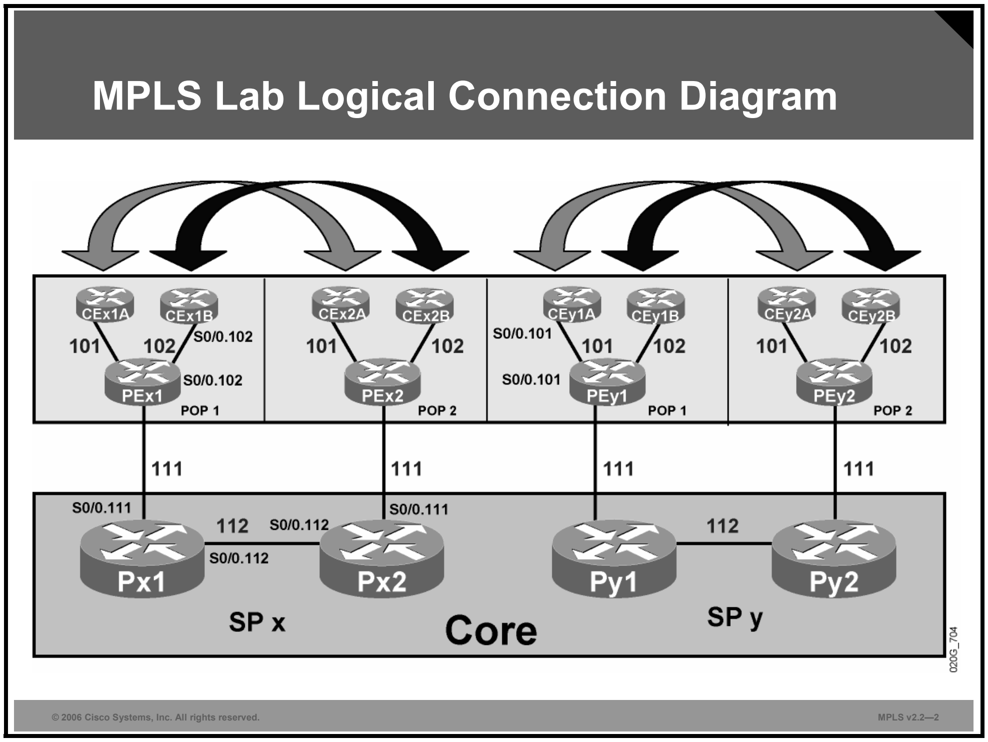

This figure represents the logical connection of two service providers. The Frame Relay DLCI information is included from the DLCI identification table. Note that the serial subinterface number matches the DLCI number.

Each SP has two P routers creating the core of the service provider network. Each P router connects to the PE router that supports the POP, which is the interface between the service provider network and the customer network. The PE routers interconnect two different customers (A and B).

Each SP is further divided into two POPs. Each POP should configure its respective left or right side of the SP. For example, SP 3 POP 1 (or POP 31) should configure P31, PE31, CE31A, and CE31B. This leaves SP 3 POP 2 (or POP 32) to configure P32, PE32, CE32A, and CE32B.

Your POP will still depend on the other POP in your SP network to complete end-to-end connectivity for customer A and customer B. Each customer has a location on each side of the POPs. An example is customer A with sites CE31A and CE32A. Site CE31A is connected to PE31 within POP 31; site CE32A is connected to the other PE32 router within POP 32.

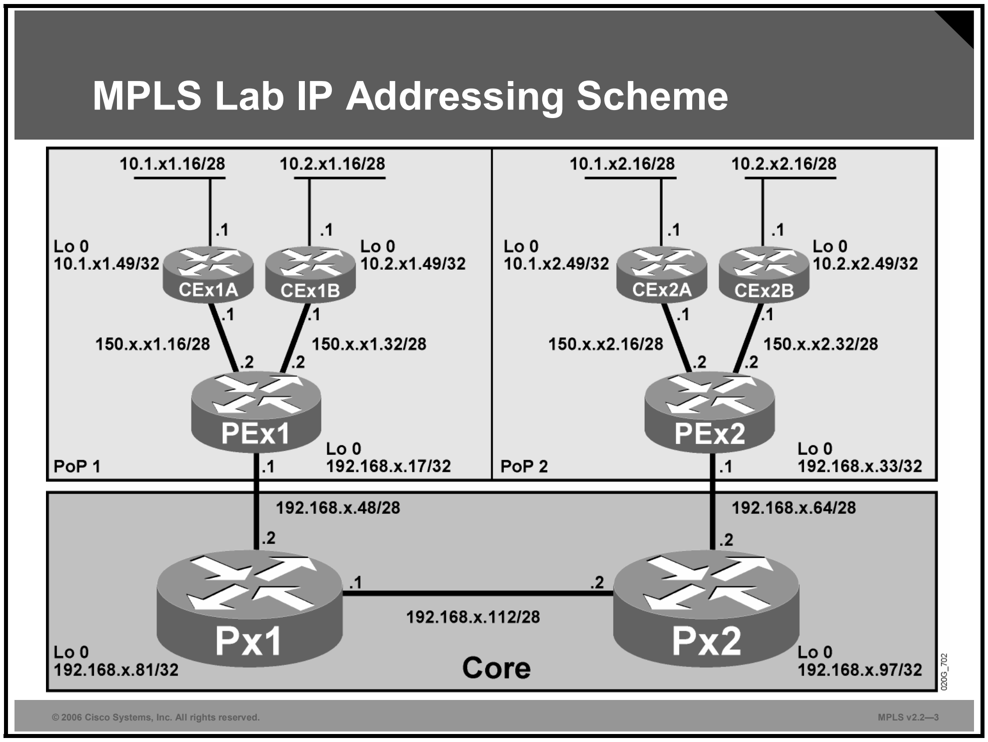

The IP addressing of routers has been performed using the allocation scheme detailed in the IP host address table. Note that *x* equals your SP number.

For all exercises, there are three distinct IP address ranges.

The 10.1.0.0 and 10.2.0.0 ranges are used to provide network addressing for the networks of customers A and B respectively. The second octet indicates the customer. The third octet of the address indicates the SP POP.

- For example, 10.**1**.**41**.16/28 is a customer A subnet on POP 41 for SP 4.

The 150.0.0.0 range is used to provide addressing for the links between the CE routers and the PE routers. The second octet of the address indicates the SP, and the third octet indicates the SP POP.

- For example, 150.**5**.**51**.16/28 is a link between a CE router (CE51) and POP 51 (or router PE51) for SP 5.

The 192.168.0.0 range is used to provide addressing for the core MPLS network of the SP. The third octet of the address indicates the SP number.

- For example, 192.168.**6**.64/28 is a link between a PE router (PE62) and a core router (P62) for SP 6.

IP Host Address

Parameter	Value
CE*x*1A (S0/0.101)	150.*x*.*x*1.17/28
CE*x*1A (loopback0)	10.1.*x*1.49/32
CE*x*1A (E0/0)	10.1.*x*1.17/28
CE*x*2A (S0/0.101)	150.*x*.*x*2.17/28
CE*x*2A (loopback0)	10.1.*x*2.49/32
CE*x*2A (E0/0)	10.1.*x*2.17/28
CE*x*1B (S0/0.102)	150.*x*.*x*1.33/28
CE*x*1B (loopback0)	10.2.*x*1.49/32
CE*x*1B (E0/0)	10.2.*x*1.17/28
CE*x*2B (S0/0.102)	150.*x*.*x*2.33/28
CE*x*2B (loopback0)	10.2.*x*2.49/32
CE*x*2B (E0/0)	10.2.*x*2.17/28
PE*x*1 (S0/0.101)	150.*x*.*x*1.18/28
PE*x*1 (S0/0.102)	150.*x*.*x*1.34/28
PE*x*1 (loopback0)	192.168.*x*.17/32
PE*x*1 (S0/0.111)	192.168.*x*.49/28
PE*x*2 (S0/0.101)	150.*x*.*x*2.18/28
PE*x*2 (S0/0.102)	150.*x*.*x*2.34/28
PE*x*2 (loopback0)	192.168.*x*.33/32
PE*x*2 (S0/0.111)	192.168.*x*.65/28
P*x*1 (S0/0.111)	192.168.*x*.50/28
P*x*1 (S0/0.112)	192.168.*x*.113/28
P*x*1 (loopback0)	192.168.*x*.81/32
P*x*2 (S0/0.111)	192.168.*x*.66/28
P*x*2 (S0/0.112)	192.168.*x*.114/28
P*x*2 (loopback0)	192.168.*x*.97/32

Note This addressing scheme has been selected for ease of use in the labs; it does not optimize the use of the address space.

Required Resources

This is the resource that is required to complete this activity:

- Cisco IOS documentation

Command List

The table describes the commands that are used in this activity.

IP, IGP, and Interface Commands

Command	Description
network *network-number* [*network-mask*]	To specify a list of networks for the EIGRP routing process, use the **network** router configuration command. To remove an entry, use the **no** form of this command.
router eigrp *as-number*	To configure the EIGRP routing process, use the **router eigrp** global configuration command. To shut down a routing process, use the **no** form of this command.
interface serial [*slot/port*]*.subinterface* **point-to-point**	To define a logical point-to-point subinterface on a physical serial interface.
encapsulation frame-relay	Enables Frame Relay encapsulation..
frame-relay interface-dlci *dlci*	Specifies the DLCI associated with its point-to-point link.
show frame-relay pvc	To display statistics about PVCs for Frame Relay interfaces, use the **show frame-relay pvc** privileged EXEC command.
show interfaces serial [*slot/port*]	To display information about a serial interface, use the **show interfaces serial** command in privileged EXEC mode. When using Frame Relay encapsulation, use the **show interfaces serial** command in EXEC mode to display information about the multicast DLCI, the DLCIs used on the interface, and the DLCI used for the LMI.
show ip protocols	To display the parameters and current state of the active routing protocol process, use the **show ip protocols** EXEC command.
show ip route [*ip-address* [*mask*] [**longer-prefixes**]] | [*protocol* [*process-id*]]	To display the current state of the routing table, use the **show ip route** EXEC command.

Task 1: Configure the Service Provider IP Interfaces

Your task is to configure Layer 2 and Layer 3 addressing and ensure that the proper interfaces are enabled.

Note	The enable password on all routers is "mpls."

Activity Procedure

Complete these steps with reference to the preceding MPLS logical connection diagram and IP addressing scheme. POP 1 and 2 of each SP should configure their respective group of routers.

Step 1 Configure and enable each P router interface, subinterface, and loopback for its appropriate DLCI and IP addressing.

Step 2 Configure and enable each PE router interface, subinterface, and loopback for its appropriate DLCI and IP addressing.

Step 3 Configure and enable each CE router interface, subinterface, and loopback for appropriate DLCI and IP addressing. Configure the Ethernet interfaces to be half-duplex with no keepalives.

Activity Verification

You have completed this task when you attain these results:

- You have pinged the remote end of each serial link from each router to verify that each link is operational.

Task 2: Configure the Service Provider IGP

Your next task is to establish the service provider IGP routing environment. This task will involve enabling the EIGRP routing protocol.

Activity Procedure

Complete these steps for POP 1 and 2 of each SP:

Step 1 On each CE router, enable the RIP version 2 (RIPv2) routing process. Advertise the 10.0.0.0 and the 150.*x*.0.0 networks. Disable the autosummarization feature of this routing protocol.

Step 2 On each P and PE router, enable the EIGRP routing process, using 1 as the AS number, and ensure that the service provider networks are configured and are being advertised by the EIGRP process. Disable the autosummarization feature of this routing protocol.

Step 3 Ensure that the other POP has completed its configuration tasks.

Activity Verification

You have completed this task when you attain these results:

- On each P and PE router, you have verified that the EIGRP router process is active.
- On each P and PE router, you have verified that the EIGRP router process is enabled on all serial interfaces.
- On each P and PE router, you have verified that the loopback interfaces of all P and PE routers are displayed in the IP routing table.
- On each P and PE router, you have verified that the 192.168.*x*.0 subnetworks of all P and PE routers are displayed in the IP routing table.
- On each PE router, you have verified that the 150.*x*.0.0 subnetworks of all P and PE routers are displayed in the IP routing table.

```
PEx1#sh 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

     192.168.x.0/24 is variably subnetted, 7 subnets, 2 masks
D       192.168.x.97/32 [90/2809856] via 192.168.x.50, 00:06:29, Serial0/0.111
D       192.168.x.112/28
           [90/2681856] via 192.168.x.50, 00:06:29, Serial0/0.111
D       192.168.x.64/28 [90/3193856] via 192.168.x.50, 00:06:29, Serial0/0.111
D       192.168.x.81/32 [90/2297856] via 192.168.x.50, 00:06:29, Serial0/0.111
D       192.168.x.33/32 [90/3321856] via 192.168.x.50, 00:06:29, Serial0/0.111
C       192.168.x.48/28 is directly connected, Serial0/0.111
C       192.168.x.17/32 is directly connected, Loopback0
     150.x.0.0/28 is subnetted, 4 subnets
C       150.x.x1.32 is directly connected, Serial0/0.102
D       150.x.x2.32 [90/3705856] via 192.168.x.50, 00:06:30, Serial0/0.111
C       150.x.x1.16 is directly connected, Serial0/0.101
D       150.x.x2.16 [90/3705856] via 192.168.x.50, 00:06:30, Serial0/0.111
PEx1#
```

```
Px2#sh 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

     192.168.x.0/24 is variably subnetted, 7 subnets, 2 masks
C       192.168.x.97/32 is directly connected, Loopback0
C       192.168.x.112/28 is directly connected, Serial0/0.112
C       192.168.x.64/28 is directly connected, Serial0/0.111
D       192.168.x.81/32
           [90/2297856] via 192.168.x.113, 00:10:55, Serial0/0.112
D       192.168.x.33/32 [90/2297856] via 192.168.x.65, 00:10:55, Serial0/0.111
D       192.168.x.48/28
           [90/2681856] via 192.168.x.113, 00:07:56, Serial0/0.112
D       192.168.x.17/32
           [90/2809856] via 192.168.x.113, 00:07:47, Serial0/0.112
     150.x.0.0/28 is subnetted, 4 subnets
D       150.x.x1.32 [90/3193856] via 192.168.x.113, 00:07:47, Serial0/0.112
D       150.x.x2.32 [90/2681856] via 192.168.x.65, 00:10:56, Serial0/0.111
D       150.x.x1.16 [90/3193856] via 192.168.x.113, 00:07:20, Serial0/0.112
D       150.x.x2.16 [90/2681856] via 192.168.x.65, 00:10:56, Serial0/0.111
Px2#
```

Lab 3-1: Establishing the Core MPLS Environment

Complete this lab activity to practice what you learned in the related module.

Activity Objective

In this activity, you will use the tasks and commands necessary to implement MPLS on frame-mode Cisco IOS platforms. After completing this activity, you will be able to meet these objectives:

- Enable LDP on your PE and P routers
- Enable and disable MPLS TTL propagation
- Configure conditional label distribution
- Remove conditional label distribution

Visual Objective

The figure illustrates what you will accomplish in this activity.

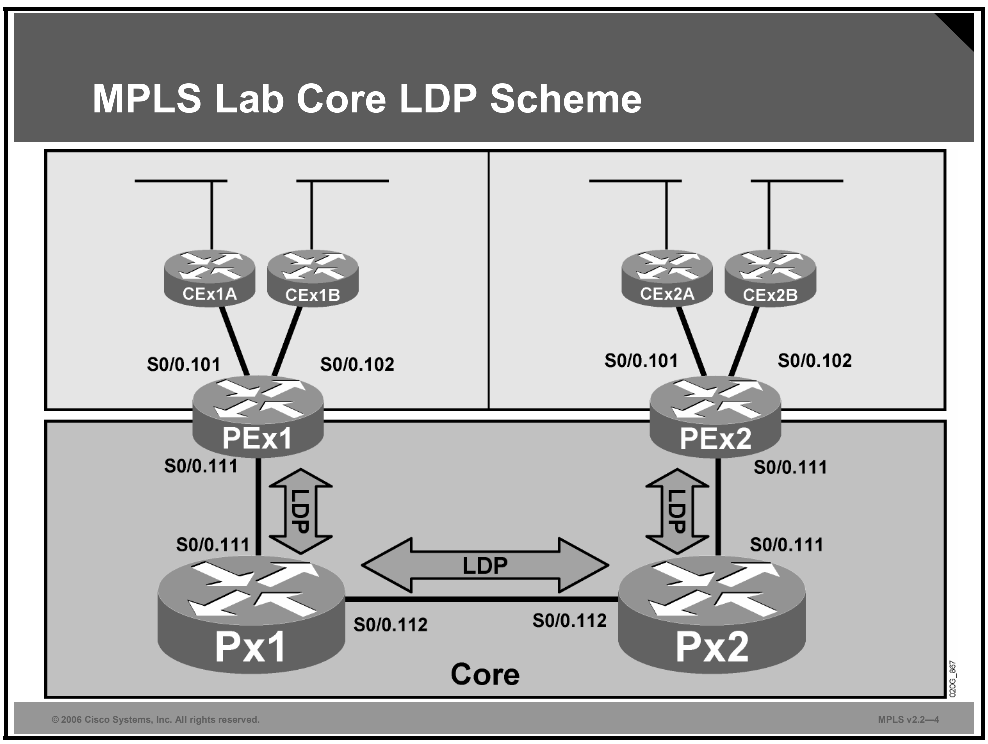

Required Resources

This is the resource that is required to complete this activity:

- Cisco IOS documentation

Command List

The table describes the commands that are used in this activity.

MPLS Commands

Command	Description
access-list *access-list-number* {**permit** | **deny**} {*type-code wild-mask* | *address mask*}	To configure the access list mechanism for filtering frames by protocol type or vendor code, use the **access-list** global configuration command. To remove the single specified entry from the access list, use the **no** form of this command.
ip cef	To enable CEF, use the **ip cef** command in global configuration mode. To disable CEF, use the **no** form of this command.
mpls ip	To enable MPLS forwarding of IPv4 packets along normally routed paths for the platform, the **mpls ip** command can be used in global configuration mode (for TE) but must be used at the interface configuration mode for LDP to become active. To disable this feature, use the **no** form of this command.
mpls ip propagate-ttl [**forwarded** | **local**]	To control the generation of the TTL field in the MPLS header when labels are first added to an IP packet, use the **mpls ip propagate-ttl** global configuration command. To use a fixed TTL value (255) for the first label of the IP packet, use the **no** form of this command.
mpls label protocol {**ldp** | **tdp** | **both** }	To specify the label distribution protocol to be used on a given interface, use the **mpls label protocol** interface configuration command. Use the **no** form of the command to disable this feature.
show mpls interfaces [*interface*] [**detail**]	To display information about one or more interfaces that have been configured for label switching, use the **show mpls interfaces** privileged EXEC command.
show mpls ldp discovery	To display the status of the LDP discovery process, use the **show mpls ldp discovery** privileged EXEC command. This command generates a list of interfaces over which the LDP discovery process is running.
show mpls ldp neighbor [*address* | *interface*] [**detail**]	To display the status of LDP sessions, issue the **show mpls ldp neighbor** privileged EXEC command.
show mpls ldp bindings [*network* {*mask* | *length*} [**longer-prefixes**]] [**local-label** *label* [- *label*]] [**remote-label** *label* [- *label*]] [**neighbor** *address*] [**local**]	To display the contents of the LIB, use the **show mpls ldp bindings** privileged EXEC command.
mpls ldp advertise-labels [**for** *prefix-access-list* [**to** *peer-access-list*]]	To control the distribution of locally assigned (incoming) labels by means of LDP, use the **mpls ldp advertise-labels** command in global configuration mode. This command is used to control which labels are advertised to which LDP neighbors. To prevent the distribution of locally assigned labels, use the **no** form of this command.

Task 1: Enable LDP on Your PE and P Routers

Your next task is to establish MPLS within the service provider routing environment. This task will involve enabling CEF and MPLS.

Activity Procedure

Complete these steps:

Step 1 On your assigned PE router, do the following:

- Enable CEF.
- Enable LDP on the subinterface that is connected to your assigned P router.

Step 2 On your assigned P router, do the following:

- Enable CEF.
- Enable LDP on the subinterface that is connected to your assigned PE router.
- Enable LDP on the subinterface that is connected to the P router of the other POP.

Note The **mpls label protocol ldp** command can be issued at the global configuration level.

Step 3 Verify that the other POP has completed its configuration.

Step 4 On your assigned PE router, determine the default TTL propagation status by using the **traceroute** command to the loopback address of the PE router of the other POP.

Note The **mpls ip command** is issued to enable MPLS on an interface, but it will be displayed in the **show running-config** command output as the **tag-switching ip** command.

Activity Verification

You have completed this task when you attain these results:

- On each of your routers, you have verified that the interfaces in question have been configured to use LDP.

```
Px1#show mpls interface
Interface              IP            Tunnel   Operational
Serial0/0.111          Yes (ldp)     No       Yes
Serial0/0.112          Yes (ldp)     No       Yes
```

- On each of your routers, you have verified that the interface is up and has established an LDP neighbor relationship.

```
Px1#show mpls ldp discovery
 Local LDP Identifier:
    192.168.1.81:0
    Discovery Sources:
    Interfaces:
        Serial0/0.111 (ldp): xmit/recv
            LDP Id: 192.168.x.17:0
        Serial0/0.112 (ldp): xmit/recv
            LDP Id: 192.168.x.97:0

Px1#show mpls ldp neighbor
    Peer LDP Ident: 192.168.x.17:0; Local LDP Ident 192.168.x.81:0
        TCP connection: 192.168.x.17.646 - 192.168.x.81.11000
        State: Oper; Msgs sent/rcvd: 20/23; Downstream
        Up time: 00:08:03
        LDP discovery sources:
          Serial0/0.111, Src IP addr: 192.168.1.49
        Addresses bound to peer LDP Ident:
          192.168.x.17    192.168.x.49    150.x.x1.18     150.x.x1.34
    Peer LDP Ident: 192.168.1.97:0; Local LDP Ident 192.168.x.81:0
        TCP connection: 192.168.x.97.11000 - 192.168.x.81.646
        State: Oper; Msgs sent/rcvd: 18/18; Downstream
        Up time: 00:06:15
        LDP discovery sources:
          Serial0/0.112, Src IP addr: 192.168.x.114
        Addresses bound to peer LDP Ident:
          192.168.x.97    192.168.x.66    192.168.x.114
```

- On each of your routers, you have verified that LDP has allocated a label for each prefix in its IP routing table.

```
PEx1#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, 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

     192.168.x.0/24 is varienably subnetted, 7 subnets, 2 masks
D       192.168.x.97/32 [90/2809856] via 192.168.x.50, 00:49:50, Serial0/0.111
D       192.168.x.112/28
           [90/2681856] via 192.168.x.50, 00:49:50, Serial0/0.111
D       192.168.x.64/28 [90/3193856] via 192.168.x.50, 00:49:50, Serial0/0.111
D       192.168.x.81/32 [90/659968] via 192.168.x.50, 00:49:50, Serial0/0.111
D       192.168.x.33/32 [90/3321856] via 192.168.1.50, 00:47:00, Serial0/0.111
C       192.168.x.48/28 is directly connected, Serial0/0.111
C       192.168.x.17/32 is directly connected, Loopback0
     150.x.0.0 is subnetted, 4 subnets
D       150.x.x2.32 [90/3705856] via 192.168.3.50, 00:50:06, Serial0/0.111
C       150.x.x1.16 is directly connected, Serial0/0.101
D       150.x.x2.16 [90/3705856] via 192.168.3.50, 00:50:06, Serial0/0.111
C       150.x.x1.32 is directly connected, Serial0/0.102
```

- On each of your routers, you have verified the LDP bindings for all prefixes of the other core routers.

```
Px1#show mpls ldp bindings
    tib entry: 150.x.x1.16/28, rev 10
        local binding:  tag: 17
        remote binding: tsr: 192.168.x.17:0, tag: imp-null
        remote binding: tsr: 192.168.x.97:0, tag: 22
  tib entry: 150.x.x1.32/28, rev 12
        local binding:  tag: 18
        remote binding: tsr: 192.168.x.17:0, tag: imp-null
        remote binding: tsr: 192.168.x.97:0, tag: 23
  tib entry: 150.x.x2.16/28, rev 21
        local binding:  tag: 22
        remote binding: tsr: 192.168.x.17:0, tag: 17
        remote binding: tsr: 192.168.x.97:0, tag: 17
  tib entry: 150.x.x2.32/28, rev 22
        local binding:  tag: 23
        remote binding: tsr: 192.168.x.17:0, tag: 16
```

```
        remote binding: tsr: 192.168.x.97:0, tag: 18
  tib entry: 192.168.x.17/32, rev 8
        local binding:  tag: 16
        remote binding: tsr: 192.168.x.17:0, tag: imp-null
        remote binding: tsr: 192.168.x.97:0, tag: 21
 tib entry: 192.168.x.33/32, rev 20
        local binding:  tag: 21
        remote binding: tsr: 192.168.x.17:0, tag: 22
        remote binding: tsr: 192.168.x.97:0, tag: 16
  tib entry: 192.168.x.48/28, rev 6
        local binding:  tag: imp-null
        remote binding: tsr: 192.168.x.17:0, tag: imp-null
        remote binding: tsr: 192.168.x.97:0, tag: 19
  tib entry: 192.168.x.64/28, rev 18
        local binding:  tag: 19
        remote binding: tsr: 192.168.x.17:0, tag: 20
        remote binding: tsr: 192.168.x.97:0, tag: imp-null
  tib entry: 192.168.x.81/32, rev 5
        local binding:  tag: imp-null
        remote binding: tsr: 192.168.x.17:0, tag: 21
        remote binding: tsr: 192.168.x.97:0, tag: 20
  tib entry: 192.168.x.97/32, rev 19
        local binding:  tag: 20
        remote binding: tsr: 192.168.x.17:0, tag: 18
        remote binding: tsr: 192.168.x.97:0, tag: imp-null
  tib entry: 192.168.x.112/28, rev 4
        local binding:  tag: imp-null
        remote binding: tsr: 192.168.x.17:0, tag: 19
        remote binding: tsr: 192.168.x.97:0, tag: imp-null
```

- On your PE router, you have determined TTL propagation status by performing a traceroute to the loopback address of the PE router of the other POP.

```
PEx2#traceroute 192.168.x.17

Type escape sequence to abort.
Tracing the route to 192.168.x.17

  1 192.168.x.49 40 msec 40 msec *
PEx2#
```


Task 2: Experiment with TTL Propagation

In this task, you will enable MPLS TTL propagation and verify the results. POP 1 will configure PE*x*1 and P*x*1. POP 2 will configure PE*x*2 and P*x*2.

Activity Procedure

Complete these steps:

Step 1 On your assigned PE router, enable MPLS TTL propagation.

Step 2 On your assigned P router, enable MPLS TTL propagation.

Step 3 Verify that the other POP has completed its configuration.

Step 4 Verify that the MPLS TTL propagation is working.

Step 5 On your assigned PE router, disable MPLS TTL propagation.

Step 6 On your assigned P router, disable MPLS TTL propagation.

Activity Verification

You have completed this task when you attain these results while MPLS TTL propagation is enabled:

- You have performed a traceroute from your PE router to the loopback address of the PE router of the other POP and verified that the results display the associated labels.

```
PEx1#trace 192.168.x.33
Type escape sequence to abort.
Tracing the route to 192.168.x.33

  1 192.168.x.50 [MPLS: Label 18 Exp 0] 196 msec 196 msec 204 msec
  2 192.168.x.114 [MPLS: Label 17 Exp 0] 100 msec 104 msec 100 msec
  3 192.168.x.65 44 msec *  40 msec

PEx2#traceroute 192.168.x.17
Type escape sequence to abort.
Tracing the route to 192.168.x.17

  1 192.168.x.66 [MPLS: Label 21 Exp 0] 184 msec 200 msec 200 msec
  2 192.168.x.113 [MPLS: Label 16 Exp 0] 100 msec 100 msec 104 msec
  3 192.168.x.49 40 msec 40 msec *
```

Note When you are troubleshooting, it may become necessary to view the core routes when doing traces. If so, it will be necessary to enable TTL propagation.

For this course, enabling TTL propagation will affect the results of some of the traces shown in the lab activity verification because additional hops and labels will be displayed.

Task 3: Configure Conditional Label Distribution

For the label binding displays that you did in Task 1, you can see that a label is assigned to every prefix that is in the IP routing table of a router. This label assignment results in wasting label space and resources to build unused LSPs. In this task, you will use conditional label advertising to restrict the distribution of labels related to the WAN interfaces in the core.

POP 1 will configure PE*x*1 and P*x*1. POP 2 will configure PE*x*2 and P*x*2.

Activity Procedure

Complete these steps:

Step 1 On your PE router, display the LSPs that are being built.

```
PEx1#show mpls forwarding-table
Local  Outgoing    Prefix            Bytes tag  Outgoing   Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
16     23          150.x.x2.32/28    0          Se0/0.111  point2point
17     22          150.x.x2.16/28    0          Se0/0.111  point2point
18     20          192.168.x.97/32   0          Se0/0.111  point2point
19     Pop tag     192.168.x.112/28  0          Se0/0.111  point2point
20     19          192.168.x.64/28   0          Se0/0.111  point2point
21     Pop tag     192.168.x.81/32   0          Se0/0.111  point2point
22     21          192.168.x.33/32   0          Se0/0.111  point2point
```

Step 2 Note that an LSP has been built to the WAN interfaces that connect the PE and P routers. This LSP will never be used because traffic will not normally terminate at this point.

Step 3 On your assigned P and PE routers, configure conditional label distribution to allow only the distribution of labels related to the core loopback addresses and the interfaces that provide direct customer support.

Step 4 Verify that the other POP has completed its configuration tasks.

Activity Verification

You have completed this task when you attain these results:

- On your PE router, you have confirmed that LSPs are not being built for the links between the PE and P routers.

```
PEx1#show mpls forwarding-table
Local  Outgoing    Prefix            Bytes tag  Outgoing   Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
16     23          150.x.x2.32/28    0          Se0/0.111  point2point
17     22          150.x.x2.16/28    0          Se0/0.111  point2point
18     20          192.168.x.97/32   0          Se0/0.111  point2point
19     Untagged    192.168.x.112/28  0          Se0/0.111  point2point
20     Untagged    192.168.x.64/28   0          Se0/0.111  point2point
21     Pop tag     192.168.x.81/32   0          Se0/0.111  point2point
22     21          192.168.x.33/32   0          Se0/0.111  point2point
```

Note An LSP is no longer built to the WAN interface that connects the other PE and P routers.

- On your P router, you have verified the LDP bindings.

```
Px1#show mpls ldp bindings
  tib entry: 150.x.x1.16/28, rev 34
        local binding:  tag: 17
        remote binding: tsr: 192.168.x.17:0, tag: imp-null
        remote binding: tsr: 192.168.x.97:0, tag: 22
  tib entry: 150.x.x1.32/28, rev 35
        local binding:  tag: 18
        remote binding: tsr: 192.168.x.17:0, tag: imp-null
        remote binding: tsr: 192.168.x.97:0, tag: 23
  tib entry: 150.x.x2.16/28, rev 36
        local binding:  tag: 22
        remote binding: tsr: 192.168.x.17:0, tag: 17
        remote binding: tsr: 192.168.x.97:0, tag: 17
  tib entry: 150.x.x2.32/28, rev 37
        local binding:  tag: 23
        remote binding: tsr: 192.168.x.17:0, tag: 16
        remote binding: tsr: 192.168.x.97:0, tag: 18
  tib entry: 192.168.x.17/32, rev 38
        local binding:  tag: 16
        remote binding: tsr: 192.168.x.17:0, tag: imp-null
        remote binding: tsr: 192.168.x.97:0, tag: 21
  tib entry: 192.168.x.33/32, rev 39
        local binding:  tag: 21
        remote binding: tsr: 192.168.x.17:0, tag: 22
        remote binding: tsr: 192.168.x.97:0, tag: 16
  tib entry: 192.168.x.48/28, rev 29
        local binding:  tag: imp-null
  tib entry: 192.168.x.64/28, rev 30
        local binding:  tag: 19
  tib entry: 192.168.x.81/32, rev 40
        local binding:  tag: imp-null
        remote binding: tsr: 192.168.x.17:0, tag: 21
        remote binding: tsr: 192.168.x.97:0, tag: 20
  tib entry: 192.168.x.97/32, rev 41
        local binding:  tag: 20
        remote binding: tsr: 192.168.x.17:0, tag: 18
        remote binding: tsr: 192.168.x.97:0, tag: imp-null
  tib entry: 192.168.x.112/28, rev 33
        local binding:  tag: imp-null
Px1#
```

Note	The prefixes assigned to the WAN interfaces connecting the other P and PE routers no longer have a remote label assigned. Further, none of the core WAN interfaces have remote labels assigned. This lessening of assignments results in a reduced label space, which saves memory resources.

Task 4: Remove Conditional Label Distribution

For the conditional label distribution displays that you did in Task 3, you can see that a label is not assigned to every prefix that is in the IP routing table of a router. In this task, you will remove conditional label advertising so that there are no restrictions on the distribution of labels related to the WAN interfaces in the core.

Note	For this simple lab environment, there are no memory constraints that would lead you to implement conditional label distribution.

POP 1 will configure PE*x*1 and P*x*1. POP 2 will configure PE*x*2 and P*x*2.

Activity Procedure

Complete these steps:

Step 1 Remove conditional label distribution and the access list that supported it.

Step 2 Verify that the other POP has completed its configuration task.

Activity Verification

You have completed this task when you attain these results:

- On your PE router, you have confirmed that all the LSPs that are being built.

```
PEx1#show mpls forwarding-table
Local  Outgoing    Prefix            Bytes tag  Outgoing   Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
16     16          192.168.x.97/32   0          Se0/0.111  point2point
17     Pop tag     192.168.x.112/28  0          Se0/0.111  point2point
18     17          192.168.x.64/28   0          Se0/0.111  point2point
19     Pop tag     192.168.x.81/32   0          Se0/0.111  point2point
20     18          192.168.x.33/32   0          Se0/0.111  point2point
21     21          150.x.x2.32/28    0          Se0/0.111  point2point
22     23          150.x.x2.16/28    0          Se0/0.111  point2point

PEx2#show mpls forwarding-table
Local  Outgoing    Prefix            Bytes tag  Outgoing   Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
16     Pop tag     192.168.5.97/32   0          Se0/0.111  point2point
17     Pop tag     192.168.5.112/28  0          Se0/0.111  point2point
18     16          192.168.5.81/32   0          Se0/0.111  point2point
19     18          192.168.5.48/28   0          Se0/0.111  point2point
20     19          192.168.5.17/32   0          Se0/0.111  point2point
21     20          150.5.x1.32/28    0          Se0/0.111  point2point
22     22          150.5.x1.16/28    0          Se0/0.111  point2point
```

Lab 5-1: Configuring Initial MPLS VPN Setup

Complete this lab activity to practice what you learned in the related module.

Activity Objective

The company that you work for is a small service provider. Your SP has been given the task of creating two simple VPNs to support two new customers (customer A and customer B) that have just signed with you.

In this activity, you will create a simple VPN for your customer. After completing this activity, you will be able to meet these objectives:

- Configure MP-BGP to establish routing between the PE routers of your POP
- Configure the VRF tables necessary to support your customer and establish your customer RIP routing using a simple VPN

Visual Objective

The figure illustrates what you will accomplish in this activity.

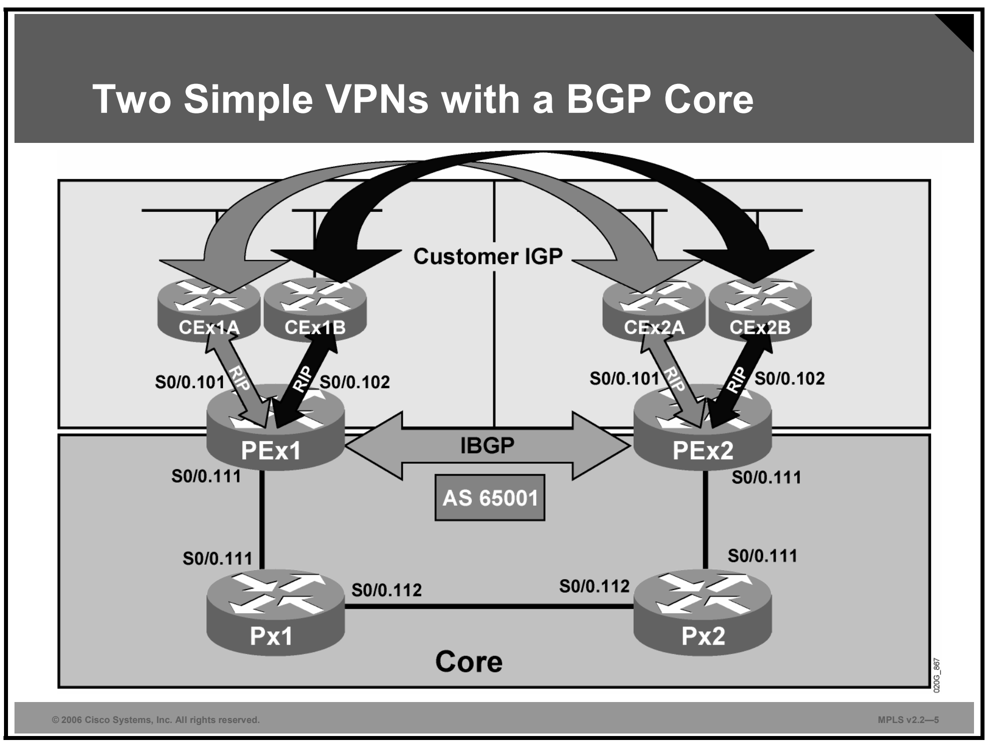

These activities rely on Lab 3-1, "Establishing the Core MPLS Environment", in which you established MPLS connectivity in your backbone.

Please verify that MPLS has been enabled on all core interfaces in your backbone and that it has not been enabled on interfaces toward the customer POP routers or other SPs.

For your reference, the figure shows the addresses in use in the network.

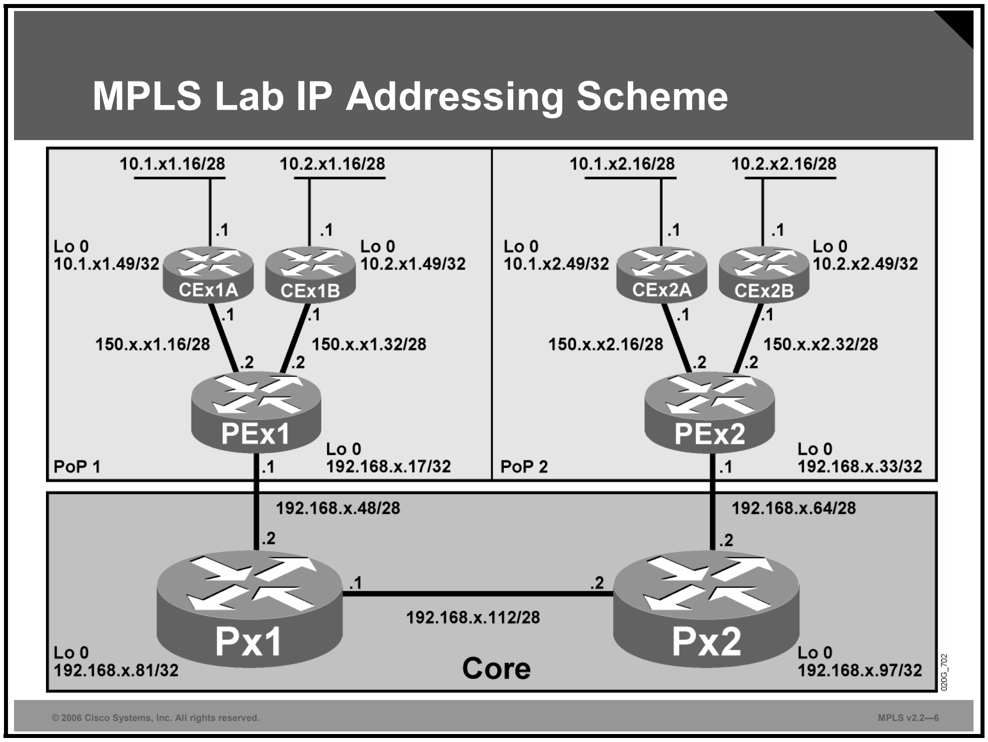

This activity contains tasks that enable you to configure your core MPLS VPN infrastructure and to establish a simple any-to-any VPN service for a customer.

You will also test various PE-CE routing options, ranging from RIP and OSPF to running BGP between the PE and the CE routers.

Required Resources

This is the resource that is required to complete this activity:

- Cisco IOS documentation

Command List

The table describes the commands that are used in this activity.

VPN-Related Commands

Command	Description
address-family ipv4 vrf *vrf-name*	Selects a per-VRF instance of a routing protocol.
address-family vpnv4	Selects VPNv4 address family configuration.
ip vrf forwarding *vrf-name*	Assigns an interface to a VRF.
ip vrf *vrf-name*	Creates a VRF table.

Command	Description
neighbor *ip-address* **activate**	Activates an exchange of routes from address family under the configuration for the specified neighbor.
neighbor *ip-address* **route-reflector-client**	Configures a route reflector client on a route reflector.
neighbor next-hop-self	To configure the router as the next hop for a BGP-speaking neighbor or peer group, use the **neighbor next-hop-self** router configuration command. To disable this feature, use the **no** form of this command.
neighbor remote-as	To add an entry to the BGP or MP-BGP neighbor table, use the **neighbor remote-as** router configuration command. To remove an entry from the table, use the **no** form of this command.
neighbor send-community	To specify that a community attribute should be sent to a BGP neighbor, use the **neighbor send-community** command in address family or router configuration mode. To remove the entry, use the **no** form of this command.
neighbor update-source	To have Cisco IOS software allow IBGP sessions to use any operational interface for TCP connections, use the **neighbor update-source** router configuration command. To restore the interface assignment to the closest interface, which is called the "best local address," use the **no** form of this command.
ping vrf *vrf-name host*	Pings a host reachable through the specified VRF.
rd *value*	Assigns an RD to a VRF.
redistribute bgp *as-number* **metric transparent**	Redistributes BGP routes into RIP with propagation of the MED into the RIP hop count.
router bgp *as-number*	Selects BGP configuration.
route-target import | export *value*	Assigns an RT to a VRF.
show ip bgp neighbor	Displays information on global BGP neighbors.
show ip bgp vpnv4 vrf *vrf-name*	Displays VPNv4 routes associated with the specified VRF.
show ip route vrf *vrf-name*	Displays an IP routing table of the specified VRF.
show ip vrf detail	Displays detailed VRF information.
show mpls forwarding vrf *vrf-name*	Displays the MPLS forwarding table for a specific VRF

Task 1: Configure MP-BGP

In this task, you will configure MP-BGP between the PE routers in your POP.

POP 1 will configure MP-BGP on PE*x*1, and POP 2 will perform the same task on PE*x*2.

Activity Procedure

Complete these steps:

Step 1 Activate the BGP process on your assigned router using AS 65001 as the AS number. Disable the autosummarization feature.

Step 2 Activate VPNv4 BGP sessions between your assigned PE router and the PE router being configured by the other POP. Disable the autosummarization feature.

Step 3 Verify that the other POP has completed its configuration tasks.

Activity Verification

You have completed this task when you attain these results:

- You have displayed the BGP neighbor information and ensured that BGP sessions have been established between the two PE routers.

```
PEx2#show ip bgp summary
BGP router identifier 192.168.x.33, local AS number 65001
BGP table version is 1, main routing table version 1
Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
192.168.x.17    4 65001       9       9        1    0    0 00:05:24        0

PEx1#show ip bgp summary
BGP router identifier 192.168.x.17, local AS number 65001
BGP table version is 1, main routing table version 1
Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd

192.168.x.33    4 65001       6       6        1    0    0 00:02:23        0

PEx1#show bgp neighbors
BGP neighbor is 192.168.x.33,  remote AS 65001, internal link
  BGP version 4, remote router ID 192.168.x.33
  BGP state = Established, up for 00:03:39
  Last read 00:00:39, 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
    IPv4 MPLS Label capability:
  Received 7 messages, 0 notifications, 0 in queue
  Sent 7 messages, 0 notifications, 0 in queue
        Default minimum time between advertisement runs is 5 seconds

       For address family: IPv4 Unicast
        BGP table version 1, neighbor version 1
        Index 1, Offset 0, Mask 0x2
        Route refresh request: received 0, sent 0
        0 accepted prefixes consume 0 bytes
        Prefix advertised 0, suppressed 0, withdrawn 0
```

```
  Connections established 1; dropped 0
  Last reset never
Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Local host: 192.168.x.17, Local port: 11022
Foreign host: 192.168.x.33, Foreign port: 179

Enqueued packets for retransmit: 0, input: 0  mis-ordered: 0 (0 bytes)

Event Timers (current time is 0xA12E784):
Timer          Starts    Wakeups            Next
Retrans             8          0             0x0
TimeWait            0          0             0x0
AckHold             7          5             0x0
SendWnd             0          0             0x0
KeepAlive           0          0             0x0
GiveUp              0          0             0x0
PmtuAger            0          0             0x0
DeadWait            0          0             0x0

iss: 1596106025  snduna: 1596106185  sndnxt: 1596106185     sndwnd:  16225
irs: 2134453172  rcvnxt: 2134453332  rcvwnd:      16225  delrcvwnd:    159

SRTT: 197 ms, RTTO: 984 ms, RTV: 787 ms, KRTT: 0 ms
minRTT: 44 ms, maxRTT: 300 ms, ACK hold: 200 ms
Flags: higher precedence, nagle

Datagrams (max data segment is 536 bytes):
Rcvd: 8 (out of order: 0), with data: 7, total data bytes: 159
Sent: 14 (retransmit: 0, fastretransmit: 0), with data: 7, total data bytes:
159
```

Task 2: Configure VRF Tables

In this task and the following task, you will establish simple VPNs for customer A and customer B. POP 1 will establish a VPN between CE*x*1A and CE*x*2A, and POP 2 will establish a VPN between CE*x*1B and CE*x*2B. Each POP is responsible for all PE router configurations related to its customer. This division of work between POPs applies to all future exercises.

Activity Procedure

Complete these steps:

Step 1 Design your VPN networks—decide on the RD and the RT numbering. Coordinate your number with the other POP.

Note The easiest numbering plan would be to use the same values for the RD and the RT. Use simple values—for example, *x*:10 for customer A and *x*:20 for customer B.

Step 2 Create VRFs on the PE routers and associate the PE-CE interfaces into the proper VRFs; use simple yet descriptive VRF names (for example, CustA and CustB).

Step 3 Your customer is using RIP as its IGP, so enable RIP for the VRF that you have created.

Step 4 Configure redistribution of RIP into BGP from the **address-family ipv4 vrf** v*rf-name* command mode.

Step 5 Configure redistribution of BGP into RIP from the **address-family ipv4 vrf** v*rf-name* command mode.

Step 6 Configure RIP metric propagation through MP-BGP by using the **redistribute bgp** *as-number* **metric transparent** command in the RIP process.

Step 7 Ensure that RIP is enabled on all of the CE routers. Make sure that all of the networks (including loopbacks) are active in the RIP process.

Activity Verification

You have completed this task when you attain these results:

- On a CE router, using the **show ip route** command, you have verified that the router is receiving all VPN routes. Also, you have verified that no routes from the other customer or the MPLS core are being received. On CE*x*1A, the printout should be similar to this example:

```
CEx1A#show ip route
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, 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, 4 subnets, 2 masks
R       10.1.x2.49/32 [120/2] via 150.x.x1.18, 00:00:14, Serial0/0.101
C       10.1.x1.49/32 is directly connected, Loopback0
R       10.1.x2.16/28 [120/2] via 150.x.x1.18, 00:00:14, Serial0/0.101
C       10.1.x1.16/28 is directly connected, Ethernet0/0
     150.x.0.0/28 is subnetted, 2 subnets
R       150.x.x2.16 [120/1] via 150.x.x1.18, 00:00:14, Serial0/0.101
C       150.x.x1.16 is directly connected, Serial0/0.101
```

- You have used **ping** and **trace** on the CE routers to verify connectivity across the VPN.

```
CEx1A#traceroute 150.x.x2.17

Type escape sequence to abort.
Tracing the route to 150.x.x2.17

  1 150.x.x1.18 12 msec 12 msec 12 msec
  2 150.x.x2.18 60 msec 60 msec 60 msec
  3 150.x.x2.17 77 msec 72 msec *

CEx1A#ping 150.x.x2.17
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.x.x2.17, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 148/148/149 ms
```

- You have verified that you have the proper configuration of your VRF tables with the **show ip vrf detail** command. Your printout should be similar to this example:

```
PEx1#show ip vrf detail
VRF CustA; default RD x:10; default VPNID <not set>
  Interfaces:
    Se0/0.101
  Connected addresses are not in global routing table
  Export VPN route-target communities
    RT:x:10
  Import VPN route-target communities
    RT:x:10
  No import route-map
  No export route-map
VRF CustB; default RD x:20; default VPNID <not set>
  Interfaces:
    Se0/0.102
  Connected addresses are not in global routing table
  Export VPN route-target communities
    RT:x:20
  Import VPN route-target communities
    RT:x:20
  No import route-map
  No export route-map
```

```
PEx2#sh ip vrf det
VRF CustA; default RD 5:10; default VPNID <not set>
  Interfaces:
    Se0/0.101
  Connected addresses are not in global routing table
  Export VPN route-target communities
    RT:x:10
  Import VPN route-target communities
    RT:x10
  No import route-map
  No export route-map
  VRF label distribution protocol: not configured
VRF CustB; default RD x:20; default VPNID <not set>
  Interfaces:
    Se0/0.102
  Connected addresses are not in global routing table
  Export VPN route-target communities
    RT:x:20
  Import VPN route-target communities
    RT:x:20
  No import route-map
  No export route-map
  VRF label distribution protocol: not configured
```

- You have checked the routing protocols running in your VRF with the **show ip protocols vrf** command. When executed on PE*x*1, your printout should be similar to this example:

```
PEx1#show ip protocols vrf CustA
Routing Protocol is "bgp 65001"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  IGP synchronization is disabled
  Automatic route summarization is disabled
  Redistributing: rip
  Maximum path: 1
  Routing Information Sources:
    Gateway         Distance      Last Update
    192.168.x.33         200      15:05:06
  Distance: external 20 internal 200 local 200

Routing Protocol is "rip"
  Sending updates every 30 seconds, next due in 26 seconds
  Invalid after 180 seconds, hold down 180, flushed after 240
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Redistributing: bgp 65001, rip
  Default version control: send version 2, receive version 2
```

```
    Interface             Send  Recv  Triggered RIP  Key-chain
    Serial0/0.101         2     2
  Maximum path: 4
  Routing for Networks:
    Interface             Send  Recv  Triggered RIP  Key-chain
    150.x.0.0
  Routing Information Sources:
    Gateway         Distance      Last Update
    150.x.x1.17          120      00:00:27
  Distance: (default is 120)
```

```
PEx1#show ip protocols vrf CustB
Routing Protocol is "bgp 65001"
 Outgoing update filter list for all interfaces is not set
 Incoming update filter list for all interfaces is not set
 IGP synchronization is disabled
 Automatic route summarization is disabled
 Redistributing: rip
 Maximum path: 1
 Routing Information Sources:
    Gateway         Distance      Last Update
    192.168.x.33         200      15:04:27
  Distance: external 20 internal 200 local 200

Routing Protocol is "rip"
  Sending updates every 30 seconds, next due in 20 seconds
  Invalid after 180 seconds, hold down 180, flushed after 240
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Redistributing: bgp 65001, rip
  Default version control: send version 2, receive version 2
    Interface             Send  Recv  Triggered RIP  Key-chain
    Serial0/0.102         2     2
  Maximum path: 4
  Routing for Networks:
    Interface             Send  Recv  Triggered RIP  Key-chain
    150.x.0.0
  Routing Information Sources:
    Gateway         Distance      Last Update
    150.x.x1.33          120      00:00:07
  Distance: (default is 120)
```

- You have verified the per-VRF routing table on the PE router with the **show ip route vrf** command. Your printout should be similar to this example:

```
PEx1#show ip route vrf CustA
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, 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, 4 subnets, 2 masks
B       10.1.x2.49/32 [200/1] via 192.168.x.33, 15:10:04
R       10.1.x1.49/32 [120/1] via 150.x.x1.17, 00:00:24, Serial0/0.101
B       10.1.x2.16/28 [200/1] via 192.168.x.33, 15:10:04
R       10.1.x1.16/28 [120/1] via 150.x.x1.17, 00:00:24, Serial0/0.101
     150.x.0.0/28 is subnetted, 2 subnets
B       150.x.x2.16 [200/0] via 192.168.x.33, 15:46:04
C       150.x.x1.16 is directly connected, Serial0/0.101

PEx1#show ip route vrf CustB
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, 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, 4 subnets, 2 masks
R     10.2.x1.49/32 [120/1] via 150.x.x1.33, 00:00:01, Serial0/0.102
B     10.2.x2.49/32 [200/1] via 192.168.x.33, 15:09:26
R     10.2.x1.16/28 [120/1] via 150.x.x1.33, 00:00:01, Serial0/0.102
B     10.2.x2.16/28 [200/1] via 192.168.x.33, 15:09:26
    150.x.0.0/28 is subnetted, 2 subnets
B     150.x.x2.32 [200/0] via 192.168.x.33, 15:46:11
C     150.x.x1.32 is directly connected, Serial0/0.102
```

- You have used the **show ip bgp vpnv4 vrf** command to display the BGP routing table associated with a VRF. The printout from the PE*x*1 router is shown here:

```
PEx1#show ip bgp vpnv4 vrf CustA
BGP table version is 47, local router ID is 192.168.x.17
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal,
              r RIB-failure
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: x:10 (default for vrf CustA)
*> 10.1.x1.16/28    150.x.x1.17              1         32768 ?
*> 10.1.x1.49/32    150.x.x1.17              1         32768 ?
*>i10.1.x2.16/28    192.168.x.33             1    100      0 ?
*>i10.1.x2.49/32    192.168.x.33             1    100      0 ?
*> 150.x.x1.16/28   0.0.0.0                  0         32768 ?
*>i150.x.x2.16/28   192.168.x.33             0    100      0 ?

PEx1#show ip bgp vpnv4 vrf CustB
BGP table version is 47, local router ID is 192.168.x.17
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal,
              r RIB-failure
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: x:20 (default for vrf CustB)
*> 10.2.x1.16/28    150.x.x1.33              1         32768 ?
*> 10.2.x1.49/32    150.x.x1.33              1         32768 ?
*>i10.2.x2.16/28    192.168.x.33             1    100      0 ?
*>i10.2.x2.49/32    192.168.x.33             1    100      0 ?
*> 150.x.x1.32/28   0.0.0.0                  0         32768 ?
*>i150.x.x2.32/28  192.168.x.33              0    100      0 ?
```

- You have used the **show mpls forwarding vrf** *vrf-name* command on the PE routers to verify the forwarding table for each VRF.

```
PEx1#show mpls forwarding-table vrf CustA
Local  Outgoing    Prefix            Bytes tag  Outgoing   Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
23     Untagged    10.1.x1.16/28[V]  0          Se0/0.101  point2point
24     Untagged    10.1.x1.49/32[V]  520        Se0/0.101  point2point
25     Aggregate   150.x.x1.16/28[V] 0

PEx1#show mpls forwarding-table vrf CustB
Local  Outgoing    Prefix            Bytes tag  Outgoing   Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
27     Untagged    10.2.x1.16/28[V]  0          Se0/0.102  point2point
28     Untagged    10.2.x1.49/32[V]  0          Se0/0.102  point2point
29     Aggregate   150.x.x1.32/28[V] 0
PEx1#
```

- You have used the **show ip route** command on the PE routers to verify that the customer routes are not in the global IP routing table.

```
PEx1#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, 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
     192.168.x.0/24 is variably subnetted, 7 subnets, 2 masks
D       192.168.x.97/32 [90/2809856] via 192.168.x.50, 19:14:54, Serial0/0.111
D       192.168.x.112/28 [90/2681856] via 192.168.x.50, 19:14:54, Serial0/0.111
D       192.168.x.64/28 [90/3193856] via 192.168.x.50, 19:14:54, Serial0/0.111
D       192.168.x.81/32 [90/2297856] via 192.168.x.50, 19:14:54, Serial0/0.111
D       192.168.x.33/32 [90/3321856] via 192.168.x.50, 19:14:54, Serial0/0.111
C       192.168.x.48/28 is directly connected, Serial0/0.111
C       192.168.x.17/32 is directly connected, Loopback0
```

- You have used the **ping** and **trace** commands on the PE routers to verify that you *cannot* reach your customer networks from global address space.

```
PEx1#ping 150.x.x1.17
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.x.x1.17, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

PEx1#ping 150.x.x1.33
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.x.x1.33, timeout is 2 seconds:
.....
```

- You have used the **ping vrf** command on the PE routers to verify that you can reach your customer networks from global address space.

```
PEx1#ping vrf CustA 150.x.x1.17
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.x.x1.17, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/31/36 ms

PEx1#ping vrf CustB 150.x.x1.33
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.x.x1.33, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/28/32 ms
```

Lab 5-2: Running EIGRP Between PE and CE Routers

Complete this lab activity to practice what you learned in the related module.

Activity Objective

Some customers use EIGRP as the routing protocol in their VPN; sometimes, EIGRP is even combined with RIP or BGP at other sites. In this activity, the customers of the service provider have decided to migrate some of their sites to EIGRP.

In this activity, you will deploy EIGRP as the PE-CE routing protocol in the VPN of your customer. After completing this activity, you will be able to meet this objective:

- Convert one of each of the customer sites to EIGRP (from RIP) and establish VPN routing using EIGRP. The other site will continue running RIP as the IGP.

Visual Objective

The figure illustrates what you will accomplish in this activity.

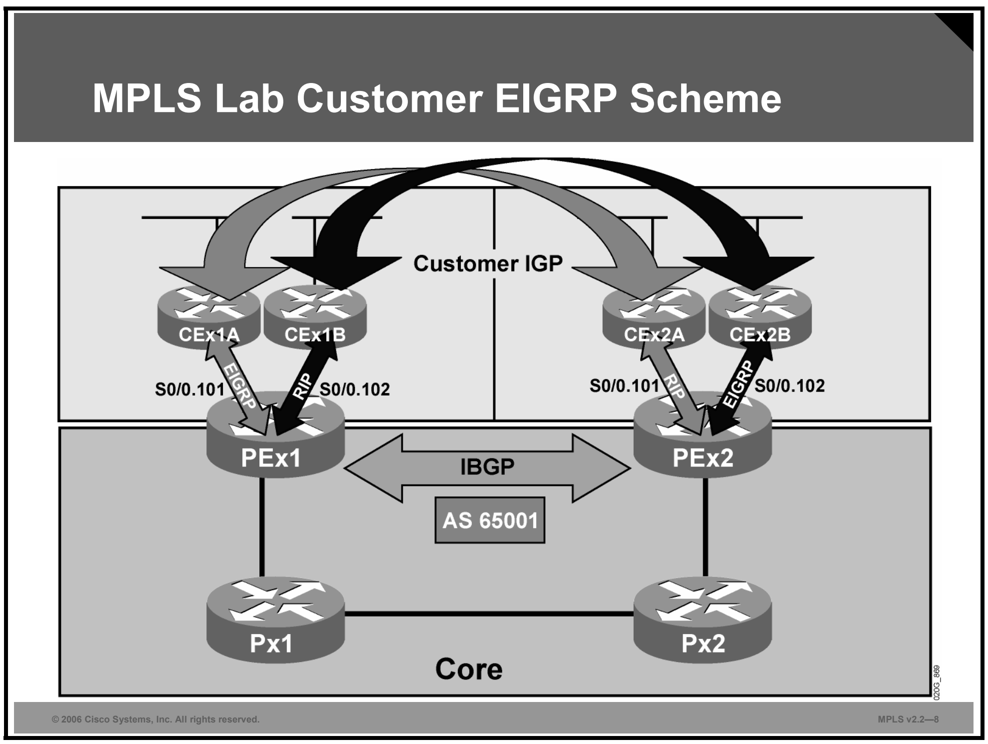

For your reference, the figure shows the addresses in use in the network.

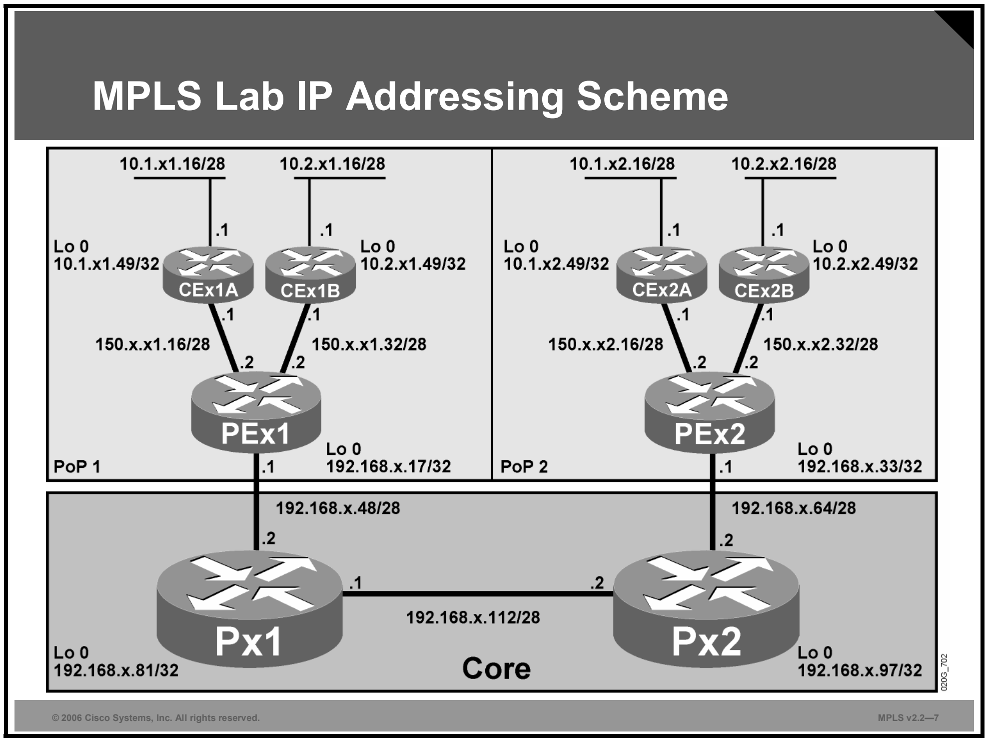

Required Resources

This is the resource that is required to complete this activity:

- Cisco IOS documentation

Command List

The table describes the commands that are used in this activity.

OSPF Commands

Command	Description
address-family ipv4 [**multicast** | **unicast** | **vrf** *vrf-name*]	Enters address family configuration mode and specifies the name of the VRF to associate with submode commands.
network *ip-address network-mask*	Specifies the network for the VRF. The network statement is used to identify which interfaces to include in EIGRP. The VRF must be configured with addresses that fall within the subnetwork range of the configured network statement.
redistribute protocol [*process-id*] {**level-1** | **level-1-2** | **level-2**} [*as-number*] [**metric** *metric-value*] [**metric-type** *type-value*] [**route-map** *map-name*] [**match** {**internal** | **external 1** | **external 2**}] [**tag** *tag-value*] [**route-map** *map-tag*] [subnets]	Redistributes BGP into the EIGRP. The AS number and metric of the BGP network are configured in this step. BGP must be redistributed into EIGRP for the CE site to accept the BGP routes that carry the EIGRP information. A metric must also be specified for the BGP network and is configured in this step.
router eigrp as-*number*	Enters router configuration mode and creates an EIGRP routing process.
show ip eigrp vrf *vrf-name* **interfaces**	Displays EIGRP interfaces that are defined under the specified VRF. If an interface is specified, only that interface is displayed. Otherwise, all interfaces on which EIGRP is running as part of the specified VRF are displayed.
show ip eigrp vrf *vrf-name* **neighbors**	Displays when VRF neighbors become active and inactive. This command can be used to help debug transport problems.
show ip eigrp vrf *vrf-name* **topology**	Displays VRF entries in the EIGRP topology table. This command can be used to determine Diffusing Update Algorithm (DUAL) states and to debug possible DUAL problems.
show ip vrf	Displays the set of defined VRFs and associated interfaces. This command is used to verify that the correct RDs are configured for the VRF.

Task 1: Enable an EIGRP VPN

In this task, your customer has decided to convert only one of its two locations from RIP to EIGRP. POP 1 will convert the customer A site, CE*x*1A, from RIP to EIGRP and establish a simple VPN.

POP 2 will convert the customer B site, CE*x*2B, from RIP to EIGRP and establish a simple VPN.

Each POP is responsible for all PE router configurations related to its customer.

Activity Procedure

Complete these steps:

Step 1 Disable RIP and configure EIGRP on one of the two routers of your customer. POP 1 will configure CE*x*1A, and POP 2 will configure CE*x*2B. Use your SP number as the AS number for EIGRP. Because both customers are connected via the same 150.*x*.0.0 network, be specific on the EIGRP statement to match the appropriate interface.

Step 2 Remove the appropriate address family from the RIP routing process on your PE router.

Step 3 Configure EIGRP on your PE router.

Step 4 On your assigned PE router, configure redistribution of EIGRP into BGP with the **address-family ipv4 vrf** *vrf-name* command. Because the source EIGRP metric is incompatible with the destination RIP metric, set the default metric to 1.

Step 5 On your assigned PE router, configure redistribution of BGP into EIRGP with the **address-family ipv4 vrf** *vrf-name* command. Disable the autosummarization feature of EIGRP.

Activity Verification

You have completed this task when you attain these results:

- You have verified that EIGRP has been activated on the proper interfaces.

```
PEx1#show ip eigrp interfaces
IP-EIGRP interfaces for process 1
                        Xmit Queue   Mean   Pacing Time   Multicast
Pending
Interface        Peers  Un/Reliable  SRTT   Un/Reliable   Flow Timer   Routes
Se0/0.111          1        0/0       600       0/15         2991           0
Lo0                0        0/0         0       0/10            0           0
```

- You have verified that EIGRP adjacencies have been established between the CE and PE routers.

```
PEx1#show ip eigrp vrf CustA neighbors
IP-EIGRP neighbors for process 4
H   Address          Interface       Hold Uptime   SRTT   RTO  Q  Seq Type
                                     (sec)         (ms)       Cnt Num
0   150.x.x1.17      Se0/0.101         14 00:02:51  340   2040  0  4

PEx2#show ip eigrp vrf CustB neighbors
IP-EIGRP neighbors for process 4
H   Address          Interface       Hold Uptime   SRTT   RTO  Q  Seq Type
                                     (sec)         (ms)       Cnt Num
0   150.x.x2.33      Se0/0.102         14 00:02:29 1050   5000  0  2
```

- You have checked the EIGRP topology database on the CE routers.

```
PEx1#show ip eigrp vrf CustA topology
IP-EIGRP Topology Table for AS(4)/ID(150.x.x1.18) Routing Table: CustA

Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

P 10.1.x2.49/32, 1 successors, FD is 281600
         via Redistributed (281600/0)
P 10.1.x1.49/32, 1 successors, FD is 2297856
         via 150.x.x1.17 (2297856/128256), Serial0/0.101
P 10.1.x2.16/28, 1 successors, FD is 281600
         via Redistributed (281600/0)
P 10.1.x1.16/28, 1 successors, FD is 2195456
         via 150.x.x1.17 (2195456/281600), Serial0/0.101
P 150.x.x2.16/28, 1 successors, FD is 281600
         via Redistributed (281600/0)
P 150.x.x1.16/28, 1 successors, FD is 2169856
         via Connected, Serial0/0.101

PEx2#show ip eigrp vrf CustB topology
IP-EIGRP Topology Table for AS(4)/ID(150.x.x2.34) Routing Table: CustB

Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

P 10.2.x1.49/32, 1 successors, FD is 281600
         via Redistributed (281600/0)
P 10.2.x1.49/32, 1 successors, FD is 2297856
         via 150.x.x2.33 (2297856/128256), Serial0/0.102
P 10.2.x1.16/28, 1 successors, FD is 281600
         via Redistributed (281600/0)
P 10.2.x2.16/28, 1 successors, FD is 2195456
         via 150.x.x2.33 (2195456/281600), Serial0/0.102
P 150.x.x2.32/28, 1 successors, FD is 2169856
         via Connected, Serial0/0.102
P 150.x.x1.32/28, 1 successors, FD is 281600
         via Redistributed (281600/0)
```

- You have verified connectivity across the VPN by using **ping** and **trace** commands on the CE routers and **ping vrf** and **trace vrf** commands on the PE routers.

```
CEx1B#ping 10.2.x2.17
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.2.x1.17, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 144/146/149 ms

CEx1A#ping 10.1.x2.17
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.x.x2.17, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 144/147/152 ms

CEx1B#trace 10.2.x2.49
Type escape sequence to abort.
Tracing the route to 150.x.x2.49
1 150.x.x1.34 16 msec 12 msec 12 msec
2 150.x.x2.34 [MPLS: Label 29 Exp 0] 64 msec 60 msec 60 msec
3 150.x.x2.33 77 msec 76 msec *

CEx1A#trace 10.1.x2.49
Type escape sequence to abort.
Tracing the route to 150.x.x2.17
1 150.x.x1.18 12 msec 12 msec 12 msec
2 150.x.x2.18 [MPLS: Label 25 Exp 0] 64 msec 60 msec 64 msec
3 150.x.x2.17 76 msec 76 msec *

PEx1#ping vrf CustA 10.1.x2.49
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.x2.49, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 116/119/120 ms

PEx1#trace vrf CustB 10.2.x2.49
Type escape sequence to abort.
Tracing the route to 10.2.x2.49

  1 150.x.x2.34 [MPLS: Label 26 Exp 0] 92 msec 88 msec 88 msec
  2 150.x.x2.33 60 msec *  60 msec
```

```
PEx2#ping vrf CustA 10.1.x1.49
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.x1.49, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/29/32 ms

PEx2#trace vrf CustA 10.1.x1.49
Type escape sequence to abort.
Tracing the route to 10.1.x1.49

  1 150.x.x1.18 [MPLS: Label 22 Exp 0] 88 msec 88 msec 88 msec
  2 150.x.x1.17 60 msec *  60 msec
```

Lab 5-3: Running OSPF Between PE and CE Routers

Complete this lab activity to practice what you learned in the related module.

Activity Objective

Some customers insist on using OSPF as the routing protocol in their VPN, sometimes even combined with RIP or BGP at other sites. In this activity, you will migrate the CE-to-PE routing protocol to OSPF. After completing this activity, you will be able to meet these objectives:

- Convert one of each of the customer sites to OSPF (from RIP) and establish VPN routing using OSPF

Visual Objective

The figure illustrates what you will accomplish in this activity.

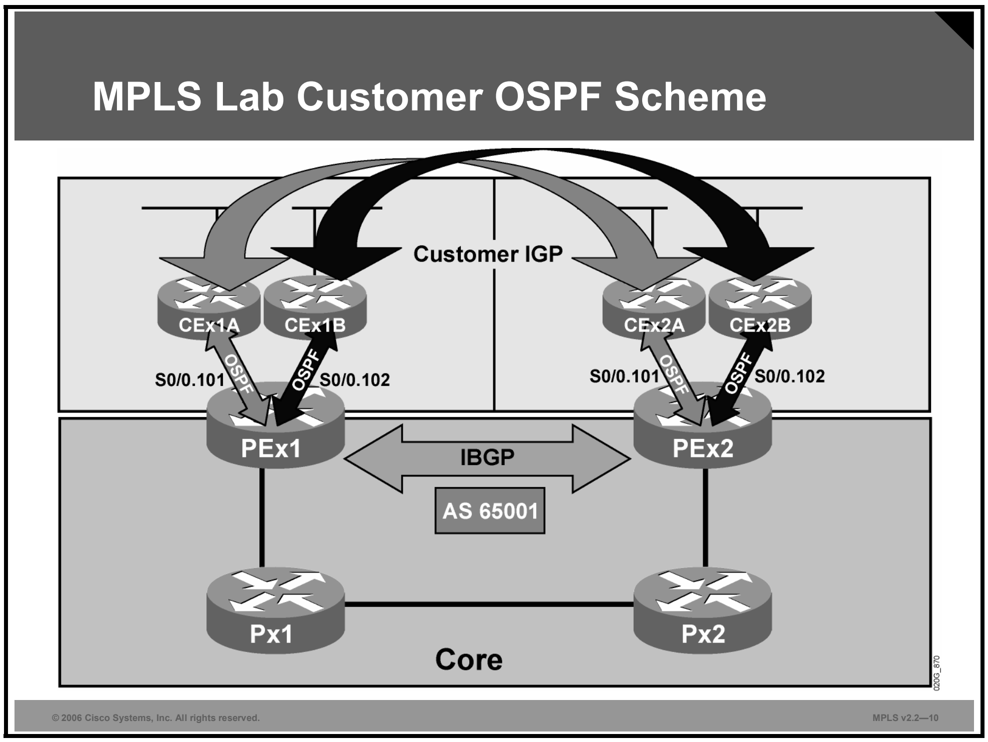

For your reference, the figure shows the addresses in use in the network.

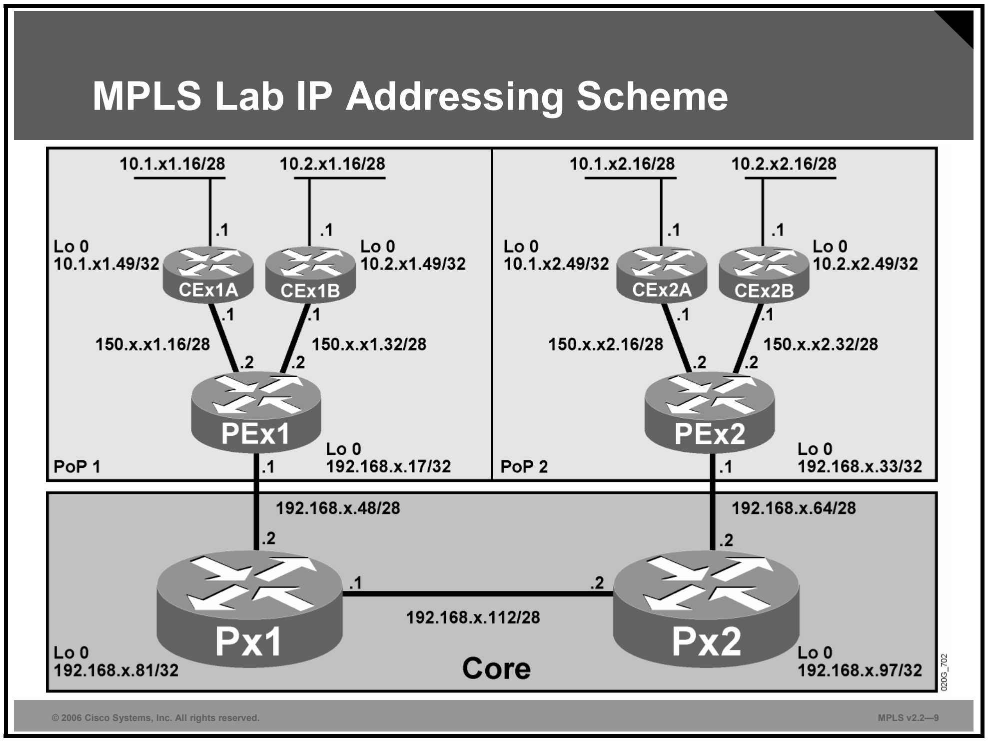

Required Resources

This is the resource that is required to complete this activity:

- Cisco IOS documentation

Command List

The table describes the commands that are used in this activity.

OSPF Commands

Command	Description
address-family ipv4 vrf *vrf-name*	Selects a per-VRF instance of a routing protocol
default-information originate always	Generates a default route into OSPF
ip vrf forwarding *vrf-name*	Assigns an interface to a VRF
ip vrf *vrf-name*	Creates a VRF table
ping vrf *vrf-name host*	Pings a host reachable through the specified VRF
rd *value*	Assigns an RD to a VRF
redistribute bgp *as-number* **subnets**	Redistributes BGP routes (including subnetwork routes) into OSPF
router bgp *as-number*	Selects BGP configuration
router ospf ***process*** **vrf** *vrf-name*	Starts an OSPF process within the specified VRF
route-target import | **export** *value*	Assigns an RT to a VRF
show ip bgp vpnv4 vrf *vrf-name*	Displays VPNv4 routes associated with the specified VRF
show ip ospf database	Displays OSPF database information
show ip route vrf *vrf-name*	Displays an IP routing table of the specified VRF
show ip vrf detail	Displays detailed VRF information
telnet host /vrf *vrf-name*	Makes a Telnet connection to a CE router connected to the specified VRF

Task 1: Configure OSPF as the PE-CE Routing Protocol

In this task, your customer has decided to have OSPF as the only IGP. This decision means that the customer sites that are running EIGRP or RIP will have to be converted to OSPF. POP 1 will convert customer A (CE*x*1A and CE*x*2A), and POP 2 will convert customer B (CE*x*1B and CE*x*2B) to establish a simple VPN.

Each POP is responsible for all PE router configurations related to its customer.

Activity Procedure

Complete these steps:

Step 1 Disable EIGRP and RIP and configure OSPF on the CE routers of your customer. Configure OSPF areas (use an OSPF process ID of 1 for CustA and a process ID of 2 for CustB) in the CE router according to the information here.

Area	Interface (or Interfaces)
Area 0	WAN interface toward PE router Loopback 0
Area 1	E0/0

Step 2 Configure OSPF (use an OSPF process ID of 1 for CustA and a process ID of 2 for CustB) in the VRFs on the PE routers using the **router ospf vrf** command. Use OSPF Area 0 on the PE-CE link.

Step 3 Configure redistribution from OSPF to MP-BGP using the **redistribute ospf** command inside the VRF address family configuration.

Step 4 Configure redistribution from MP-BGP to OSPF using the **redistribute bgp subnets** command in the OSPF router configuration.

Activity Verification

You have completed this task when you attain these results:

- You have verified the OSPF adjacency on PE*x*1 and PE*x*2 routers using the **show ip ospf neighbor** command.

```
PEx1#show ip ospf neighbor
Neighbor ID     Pri   State         Dead Time   Address         Interface
10.1.x1.49        0   FULL/  -      00:00:36    150.x.x1.17     Serial0/0.101
10.2.x1.49        0   FULL/  -      00:00:37    150.x.x1.33     Serial0/0.102

PEx2#show ip ospf neighbor
Neighbor ID     Pri   State         Dead Time   Address         Interface
10.2.x2.49        0   FULL/  -      00:00:30    150.x.x2.33     Serial0/0.102
10.1.x2.49        0   FULL/  -      00:00:39    150.x.x2.17     Serial0/0.101
```

- You have checked the OSPF topology database on CE*x*1A and CE*x*2B. You should see router link states (resulting from OSPF connectivity between the PE and the CE routers) and type 5 external link states. A sample printout from CE*x*1A is shown here:

```
CEx1A#show ip ospf database

           OSPF Router with ID (10.1.11.49) (Process ID 1)

                Router Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum Link count
10.1.x1.49      10.1.x1.49      1744        0x20000005 0x007C30 3
150.x.x1.18     150.x.x1.18     216         0x20000004 0x000E87 2


                Summary Net Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum
10.1.x1.16      10.1.x1.49      1744        0x20000002 0x0012C1
10.1.x2.16      150.x.x1.18     1186        0x20000001 0x00CDD7
10.1.x2.49      150.x.x1.18     1186        0x20000001 0x0082FB
150.x.x2.16     150.x.x1.18     1186        0x20000001 0x00CD94


                Router Link States (Area 1)

Link ID         ADV Router      Age         Seq#       Checksum Link count
10.1.x1.49      10.1.x1.49      1744        0x20000002 0x00532E 1


                Summary Net Link States (Area 1)

Link ID         ADV Router      Age         Seq#       Checksum
10.1.x1.49      10.1.x1.49      1744        0x20000002 0x00C6E5
10.1.x2.16      10.1.x1.49      1294        0x20000001 0x000E45
10.1.x2.49      10.1.x1.49      1294        0x20000001 0x00C269
150.x.x1.16     10.1.x1.49      1853        0x20000002 0x000D04
150.x.x2.16     10.1.x1.49      1294        0x20000001 0x000E02


                Summary ASB Link States (Area 1)

Link ID         ADV Router      Age         Seq#       Checksum
150.x.x1.18     10.1.x1.49      332         0x20000002 0x0045B9
```

- You have checked the IP routing table on CE*x*1A and noted the OSPF interarea (IA) routes in the routing table.

```
CEx1A#show ip route
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, 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, 4 subnets, 2 masks
C       10.1.x1.16/28 is directly connected, Ethernet0/0
O IA    10.1.x2.16/28 [110/138] via 150.x.x1.18, 00:32:41, Serial2/0.101
C       10.1.x1.49/32 is directly connected, Loopback0
O IA    10.1.x2.49/32 [110/129] via 150.x.x1.18, 00:32:41, Serial2/0.101
     150.x.0.0/28 is subnetted, 2 subnets
O IA    150.x.x2.16 [110/65] via 150.x.x1.18, 00:32:41, Serial2/0.101
C       150.x.x1.16 is directly connected, Serial2/0.101
```

- You have verified connectivity across the VPN by using **ping** and **trace** commands on the CE routers and **ping vrf** and **trace vrf** commands on the PE routers. These are just a few examples.

```
CEx1A#ping 10.1.x2.49
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.x2.49, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 148/148/149 ms

PEx1#ping vrf CustA 10.1.x2.49
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.x2.49, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 118/122/128 ms

PEx1#ping vrf CustB 10.2.x2.49
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.2.x2.49, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 116/121/132 ms
```

```
PEx1#trace vrf CustA 10.1.x2.49
Type escape sequence to abort.
Tracing the route to 10.1.x2.49

  1 150.x.x2.18 [MPLS: Label 16 Exp 0] 88 msec 92 msec 88 msec
  2 150.x.x2.17 60 msec *  60 msec

PEx1#trace vrf CustB 10.2.x2.49
Type escape sequence to abort.
Tracing the route to 10.2.x2.49

  1 150.x.x2.34 [MPLS: Label 24 Exp 0] 88 msec 88 msec 88 msec
  2 150.x.x2.33 60 msec *  60 msec
```

Lab 5-4: Running BGP Between PE and CE Routers

Complete this lab activity to practice what you learned in the related module.

Activity Objective

Your customer has indicated that it wants to have a backup link for a selected site for redundancy. This addition will produce a multihomed environment. As a result, it is necessary to use BGP as the CE-to-PE routing protocol. The provider has decided to implement this conversion in phases. The existing links will be converted to BGP, and then the backup links will be added and activated.

In this activity, you will convert the CE-to-PE routing protocol of your customer to BGP. After completing this activity, you will be able to meet these objectives:

- Enable EBGP as the CE-to-PE link routing protocol
- Enable a backup link
- Configure BGP to control the selection of primary and backup links

Visual Objective

The figure illustrates what you will accomplish in this activity.

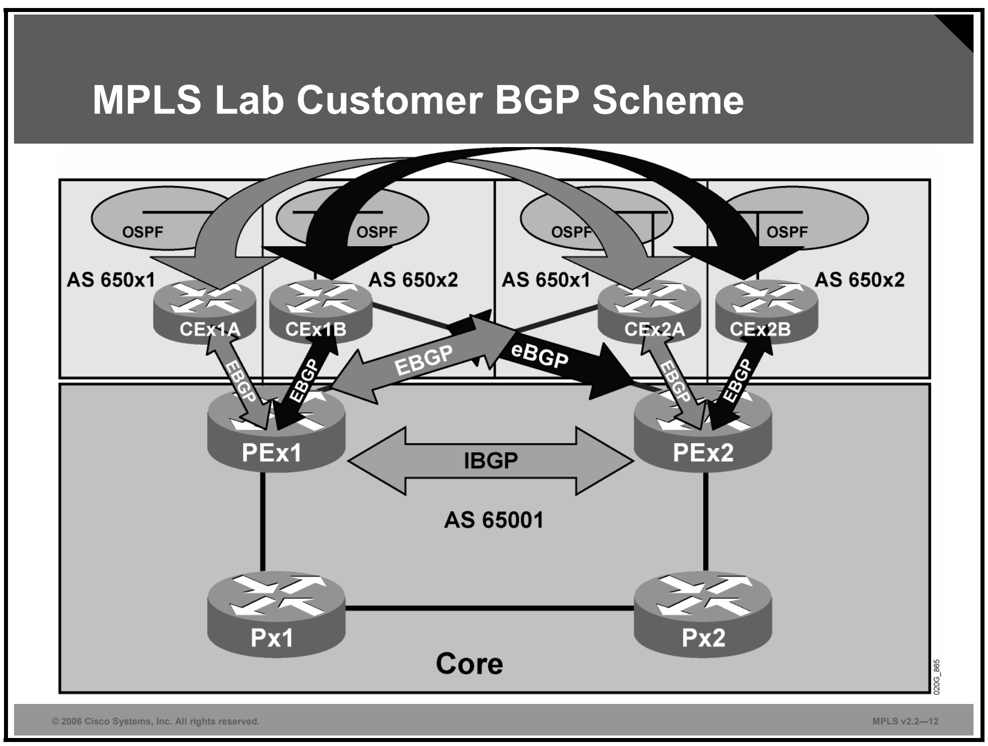

For your reference, the figure shows the addresses in use in the network.

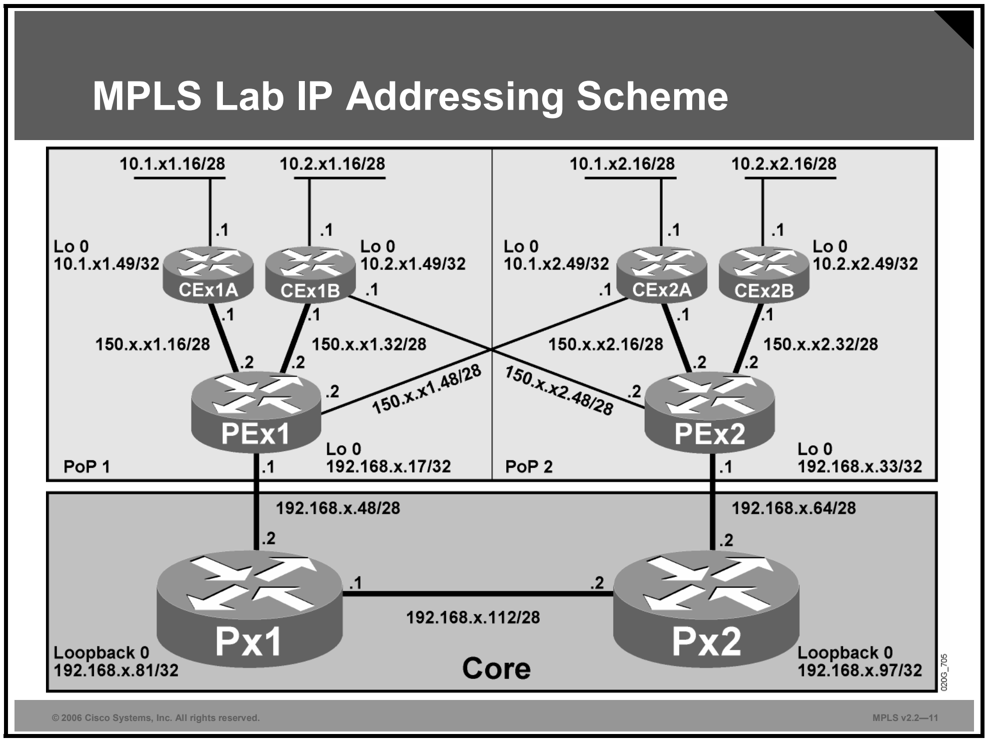

Required Resources

This is the resource that is required to complete this activity:

- Cisco IOS documentation

Command List

The table describes the commands that are used in this activity.

BGP Commands

Command	Description
address-family ipv4 vrf *vrf-name*	Selects a per-VRF instance of a routing protocol.
ip vrf forwarding *vrf-name*	Assigns an interface to a VRF.
ip vrf *vrf-name*	Creates a VRF table.
neighbor *ip-address* **as-override**	To configure a PE router to override the AS number of a site with the AS number of a provider, use the **neighbor as-override** command in router configuration mode. To remove VPNv4 prefixes from a specified router, use the **no** form of this command.
neighbor *ip-address* **route-map** *name* **in** | **out**	Applies a route map to BGP updates received from or sent to the specified neighbor.
no neighbor *ip-address* **shutdown**	Enables a BGP neighbor previously disabled with the **neighbor shutdown** command.
ping vrf *vrf-name host*	Pings a host reachable through the specified VRF.
rd *value*	Assigns an RD to a VRF.
route-map *name* **permit** *seq*	Creates an entry in a route map.
router bgp *as-number*	Selects BGP configuration.
route-target import | **export** *value*	Assigns an RT to a VRF.
set metric *value*	Sets the BGP MED attribute in a route map.
show ip bgp vpnv4 vrf *vrf-name*	Displays VPNv4 routes associated with the specified VRF.
show ip route vrf *vrf-name*	Displays an IP routing table of the specified VRF.
telnet host /vrf *vrf-name*	Makes a Telnet connection to a CE router connected to the specified VRF.

Task 1: Configure BGP as the PE-CE Routing Protocol

In this task, you will make BGP the routing protocol between the PE router and your customer routers. OSPF will remain the customer IGP. You will need to redistribute from BGP to OSPF and from OSPF to BGP on the routers of your customer. You will establish simple VPNs for customer A and customer B. POP 1 will convert customer A (CE*x*1A and CE*x*2A), and POP 2 will convert customer B (CE*x*1B and CE*x*2B) to establish a simple VPN. Each POP is responsible for all PE router configurations related to its customer.

Activity Procedure

Complete these steps:

Step 1 Activate the BGP routing process on the CE routers of your customer using AS650*x*1 for customer A and AS 650*x*2 for customer B. Disable the autosummarization BGP feature.

Step 2 Remove OSPF on the associated PE router and activate the BGP neighbor relationship between each CE router and its associated PE router.

Step 3 Because both customers use the same AS number at all their sites, you will need to enable the AS-override feature on the PE routers.

Activity Verification

You have completed this task when you attain these results:

- You have checked BGP connectivity with the **show ip bgp summary** command on the CE routers.

```
CEx1A#show ip bgp summary
BGP router identifier 10.1.x1.49, local AS number 650x1
BGP table version is 10, main routing table version 10
6 network entries using 582 bytes of memory
6 path entries using 216 bytes of memory
1 BGP AS-PATH entries using 24 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 942 total bytes of memory
BGP activity 6/0 prefixes, 9/3 paths, scan interval 60 secs

Neighbor     V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
150.x.x1.18  4 65001     617     618       10    0    0 09:50:35        3

CEx1A#show ip bgp
BGP table version is 7, local router ID is 10.1.x1.49
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
*> 10.1.x1.16/28    0.0.0.0                  0         32768 ?
*> 10.1.x1.49/32    0.0.0.0                  0         32768 ?
*> 10.1.x2.16/28    150.x.x1.18                            0 65001 65001 ?
*> 10.1.x2.49/32    150.x.x1.18                            0 65001 65001 ?
*> 150.x.x1.16/28   0.0.0.0                  0         32768 ?
*> 150.x.x2.16/28   150.x.x1.18                            0 65001 65001 ?
```

```
PEx1#show ip bgp vpn all
BGP table version is 63, local router ID is 192.168.x.17
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: x:10 (default for vrf CustA)
*> 10.1.x1.16/28    150.x.x1.17              0             0 650x1 ?
*> 10.1.x1.49/32    150.x.x1.17              0             0 650x1 ?
*>i10.1.x2.16/28    192.168.x.33             0    100      0 650x1 ?
*>i10.1.x2.49/32    192.168.x.33             0    100      0 650x1 ?
r> 150.x.x1.16/28   150.x.x1.17              0             0 650x1 ?
*>i150.x.x2.16/28   192.168.x.33             0    100      0 650x1 ?
Route Distinguisher: x:20 (default for vrf CustB)
*> 10.2.x1.16/28    150.x.x1.33              0             0 650x2 ?
*> 10.2.x1.49/32    150.x.x1.33              0             0 650x2 ?
*>i10.2.x2.16/28    192.168.x.33             0    100      0 650x2 ?
*>i10.2.x2.49/32    192.168.x.33             0    100      0 650x2 ?
r> 150.x.x1.32/28   150.x.x1.33              0             0 650x2 ?
*>i150.x.x2.32/28   192.168.x.33             0    100      0 650x2 ?
```

Task 2: Configure the Backup PE-CE Link

In this task, you will enable the backup links on the PE routers. POP 1 will establish the link between its PE*x*1 router and the CE*x*2A router, and POP 2 will establish the link between its PE*x*2 router and the CE*x*1B router. Ensure that the interface is added to the proper VRF and that BGP is activated.

Activity Procedure

Complete these steps:

Step 1 Configure an additional subinterface on the existing serial interfaces on your PE and CE routers. Use the DLCI number as the subinterface number.

Step 2 Add the backup link to the appropriate VRF.

Which VRF is interface Se0/0.113 from CE*x*1B added to?

____

Which VRF is interface Se0/0.113 from CE*x*2A added to?

____

Step 3 Configure IP addresses and DLCIs on this interface using the parameters in the table.

Backup Link Configuration Parameters

Source Router	IP Address	DLCI	Destination Router	IP Address	DLCI
CE*x*2A	150.*x*.*x*1.49/28	113	PE*x*1	150.*x*.*x*1.50/28	113
CE*x*1B	150.*x*.*x*2.49/28	113	PE*x*2	150*x*.*x*2.50/28	113

Step 4 Activate the BGP neighbor relationship between your CE router and the appropriate PE router.

Activity Verification

You have completed this task when you attain these results:

- You have verified point-to-point connectivity over the new subinterface.

```
CEx1B#ping 150.x.x2.50
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.x.x2.50, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/28/32 ms

PEx2#ping vrf CustB 150.x.x2.49
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.x.x2.49, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/56/60 ms

CEx2A#ping 150.x.x1.50
 Sending 5, 100-byte ICMP Echos to 150.x.x1.50, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/29/32 ms

PEx1#ping vrf CustA 150.x.x1.49
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.x.x1.49, timeout is 2 seconds:
!!!!!
```

- You have checked BGP connectivity with the **show ip bgp summary** command on the CE routers.

```
CEx2A#show ip bgp summary
BGP router identifier 10.1.x2.49, local AS number 650x2
BGP table version is 10, main routing table version 10
7 network entries and 9 paths using 1197 bytes of memory
10 path entries using 360 bytes of memory
1 BGP AS-PATH entries using 24 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 1183 total bytes of memory
BGP activity 7/0 prefixes, 16/6 paths, scan interval 60 secs

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down
State/PfxRcd
150.x.x1.50     4 65001     606     607       10    0    0 00:01:29        2
150.x.x2.18     4 65001     617     618       10    0    0 09:50:35        3

CEx2A#show ip bgp
BGP table version is 17, local router ID is 10.1.x2.49
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
*  10.1.x1.16/28    150.x.x2.18                            0 65001 65001 ?
*>                  150.x.x1.50                            0 65001 65001 ?
*  10.1.x1.49/32    150.x.x2.18                            0 65001 65001 ?
*>                  150.x.x1.50                            0 65001 65001 ?
*> 10.1.x2.16/28    0.0.0.0                  0         32768 ?
*> 10.1.x2.49/32    0.0.0.0                  0         32768 ?
*  150.x.x1.16/28   150.x.x2.18                            0 65001 65001 ?
*>                  150.x.x1.50                            0 65001 65001 ?
*> 150.x.x1.48/28   0.0.0.0                  0         32768 ?
*> 150.x.x2.16/28   0.0.0.0                  0         32768 ?
```

```
PEx1#show ip bgp vpn all
BGP table version is 36, local router ID is 192.168.x.17
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: x:10 (default for vrf CustA)
*> 10.1.x1.16/28    150.x.x1.17              0             0 650x1 ?
*> 10.1.x1.49/32    150.x.x1.17              0             0 650x1 ?
*> 10.1.x2.16/28    150.x.x1.49              0             0 650x1 ?
* i                 192.168.x.33             0    100      0 650x1 ?
*> 10.1.x2.49/32    150.x.x1.49              0             0 650x1 ?
* i                 192.168.x.33             0    100      0 650x1 ?
r> 150.x.x1.16/28   150.x.x1.17              0             0 650x1 ?
r> 150.x.x1.48/28   150.x.x1.49              0             0 650x1 ?
r i                 192.168.x.33             0    100      0 650x1 ?
*> 150.x.x2.16/28   150.x.x1.49              0             0 650x1 ?
* i                 192.168.1.33             0    100      0 650x1 ?
Route Distinguisher: x:20 (default for vrf CustB)
* i10.2.x1.16/28    192.168.x.33             0    100      0 650x2 ?
*>                  150.x.x1.33              0             0 650x2 ?
* i10.2.x1.49/32    192.168.x.33             0    100      0 650x2 ?
*>                  150.x.x1.33              0             0 650x2 ?
*>i10.2.x2.16/28    192.168.x.33             0    100      0 650x2 ?
*>i10.2.x2.49/32    192.168.x.33             0    100      0 650x2 ?
r i150.x.x1.32/28   192.168.x.33             0    100      0 650x2 ?
r>                  150.x.x1.33              0             0 650x2 ?
*>i150.x.x2.32/28   192.168.x.33             0    100      0 650x2 ?
* i150.x.x2.48/28   192.168.x.33             0    100      0 650x2 ?
*>                  150.x.x1.33              0             0 650x2 ?
```

```
PEx2#show ip bgp vpn all
BGP table version is 130, local router ID is 192.168.x.33
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: x:10 (default for vrf CustA)
*>i10.1.x1.16/28    192.168.x.17             0    100      0 650x1 ?
*>i10.1.x1.49/32    192.168.x.17             0    100      0 650x1 ?
* i10.1.x2.16/28    192.168.x.17             0    100      0 650x1 ?
*>                  150.x.x2.17              0             0 650x1 ?
* i10.1.x2.49/32    192.168.x.17             0    100      0 650x1 ?
*>                  150.x.x2.17              0             0 650x1 ?
*>i150.x.x1.16/28   192.168.x.17             0    100      0 650x1 ?
* i150.x.x1.48/28   192.168.x.17             0    100      0 650x1 ?
*>                  150.x.x2.17              0             0 650x1 ?
r i150.x.x2.16/28   192.168.x.17             0    100      0 650x1 ?
r>                  150.x.x2.17              0             0 650x1 ?
Route Distinguisher: x:20 (default for vrf CustB)
*> 10.2.x1.16/28    150.x.x2.49              0             0 650x2 ?
* i                 192.168.x.17             0    100      0 650x2 ?
*> 10.2.x1.49/32    150.x.x2.49              0             0 650x2 ?
* i                 192.168.x.17             0    100      0 650x2 ?
*> 10.2.x2.16/28    150.x.x2.33              0             0 650x2 ?
*> 10.2.x2.49/32    150.x.x2.33              0             0 650x2 ?
*> 150.x.x1.32/28   150.x.x2.49              0             0 650x2 ?
* i                 192.168.x.17             0    100      0 650x2 ?
r> 150.x.x2.32/28   150.x.x2.33              0             0 650x2 ?
r> 150.x.x2.48/28   150.x.x2.49              0             0 650x2 ?
r i                 192.168.x.17             0    100      0 650x2 ?
```

Task 3: Select the Primary and Backup Link with BGP

It may be necessary to control the BGP selection of the link to establish a primary backup relationship. In this task, you will use the local preference and MED attributes to control link selection. In this implementation, the new link bypasses the MPLS core. However, because it is a high-cost link, it should be considered only as the backup link; the link through the MPLS core is to be used as the primary link.

Activity Procedure

Complete these steps:

Step 1 Use the BGP local preference on the CE router to select the link to its local PE router (through the MPLS core) as the primary link and the link to the remote PE router (bypass link) as the backup link. Use a lower local preference on the preferred path.

Step 2 Set the MED in outgoing routing updates from your CE router to make sure that the PE routers prefer the link through the MPLS core before using the backup link.

Activity Verification

You have completed this task when you attain these results:

- You may have had to issue a **clear ip route** or **clear ip bgp *** command on the CE router to propagate routes with the new parameters.
- You have verified that the primary link (the link to your local PE router) is being used. Use the **show ip bgp** command to verify this. Make sure that the routes received from the primary link are always selected as the best routes.

```
CEx1B#show ip bgp
BGP table version is 8, local router ID is 10.2.x1.49
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
*  10.2.x1.16/28    150.x.x2.50                    50      0 65001 65001 ?
*>                  0.0.0.0                  0         32768 ?
*> 10.2.x1.49/32    0.0.0.0                  0         32768 ?
*  10.2.x2.16/28    150.x.x2.50                    50      0 65001 65001 ?
*>                  150.x.x1.34                            0 65001 65001 ?
*  10.2.x2.49/32    150.x.x2.50                    50      0 65001 65001 ?
*>                  150.x.x1.34                            0 65001 65001 ?
*  150.x.x1.32/28   150.x.x2.50                    50      0 65001 65001 ?
*>                  0.0.0.0                  0         32768 ?
*  150.x.x2.32/28   150.x.x2.50                    50      0 65001 65001 ?
*>                  150.x.x1.34                            0 65001 65001 ?
*  150.x.x2.48/28   150.x.x2.50                    50      0 65001 65001 ?
*>                  0.0.0.0                  0         32768 ?
```

- You have verified the proper setting of the MED by using the **show ip bgp vpnv4 vrf** command on the PE routers. Make sure that the PE routers select routes coming from the primary link as the best routes.

```
PEx2#show ip bgp vpnv4 all
BGP table version is 30, local router ID is 192.168.x.33
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: x:10 (default for vrf CustA)
*>i10.1.x1.16/28    192.168.x.17             0    100      0 650x1 ?
*>i10.1.x1.49/32    192.168.x.17             0    100      0 650x1 ?
*> 10.1.x2.16/28    150.x.x2.17              0             0 650x1 ?
*> 10.1.x2.49/32    150.x.x2.17              0             0 650x1 ?
*>i150.x.x1.16/28   192.168.x.17             0    100      0 650x1 ?
*> 150.x.x1.48/28   150.x.x2.17              0             0 650x1 ?
r> 150.x.x2.16/28   150.x.x2.17              0             0 650x1 ?
Route Distinguisher: x:20 (default for vrf CustB)
*>i10.2.x1.16/28    192.168.x.17             0    100      0 650x2 ?
*                   150.x.x2.49            200             0 650x2 ?
*  10.2.x1.49/28    150.x.x2.49            200             0 650x2 ?
*>i                 192.168.x.17             0    100      0 650x2 ?
*> 10.2.x2.16/28    150.x.x2.33              0             0 650x2 ?
*> 10.2.x2.49/32    150.x.x2.33              0             0 650x2 ?
*>i150.x.x1.32/28   192.168.x.17             0    100      0 650x2 ?
*                   150.x.x2.49            200             0 650x2 ?
r> 150.x.x2.32/28   150.x.x2.33              0             0 650x2 ?
r>i150.x.x2.48/28   192.168.x.17             0    100      0 650x2 ?
r                   150.x.x2.49            200             0 650x2 ?
```

- You have shut down the link from the local PE router to the dual-connected CE router. (This is interface Se0/0.102 on PE*x*1, and Se0/0.101 on PE*x*2.)

- You have verified that the backup link (the link to your remote PE router) is being used. Use the **show ip bgp** command to verify this after BGP converges.

```
CEx1B#show ip bgp
BGP table version is 7, local router ID is 10.2.x1.49
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
*> 10.2.x1.16/28    0.0.0.0                  0         32768 ?
*> 10.2.x1.49/32    0.0.0.0                  0         32768 ?
*> 10.2.x2.16/28    150.x.x2.50                     50      0 65001 65001 ?
*> 10.2.x2.49/32    150.x.x2.50                     50      0 65001 65001 ?
*> 150.x.x2.32/28   150.x.x2.50                     50      0 65001 65001 ?
*> 150.x.x2.48/28   0.0.0.0                  0         32768 ?

CEx2A#sh ip bgp
BGP table version is 9, local router ID is 10.1.x2.49
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
*> 10.1.x1.16/28    150.x.x1.50                     50      0 65001 65001 ?
*> 10.1.x1.49/32    150.x.x1.50                     50      0 65001 65001 ?
*> 10.1.x2.16/28    0.0.0.0                  0         32768 ?
*> 10.1.x2.49/32    0.0.0.0                  0         32768 ?
*> 150.x.x1.16/28   150.x.x1.50                     50      0 65001 65001 ?
*> 150.x.x1.48/28   0.0.0.0                  0         32768 ?
```

- You have re-enabled the subinterface.

- After the BGP session is established with the local PE router, you have verified that the local link is shown as the preferred link for traffic. Use the **show ip bgp** command to verify this.

```
CEx1B#show ip bgp

BGP table version is 8, local router ID is 10.2.x1.49

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

*  10.2.x1.16/28    150.x.x2.50                     50      0 65001 65001 ?

*>                  0.0.0.0                  0         32768 ?

*  10.2.x1.49/32    0.0.0.0                  0         32768 ?

*  10.2.x2.16/28    150.x.x2.50                     50      0 65001 65001 ?

*>                  150.x.x1.34                             0 65001 65001 ?

*  10.2.x2.49/32    150.x.x2.50                     50      0 65001 65001 ?

*>                  150.x.x1.34                             0 65001 65001 ?

*  150.x.x1.32/28   150.x.x2.50                     50      0 65001 65001 ?

*>                  0.0.0.0                  0         32768 ?

*  150.x.x2.32/28   150.x.x2.50                     50      0 65001 65001 ?

*>                  150.x.x1.34                             0 65001 65001 ?

*  150.x.x2.48/28   150.x.x2.50                     50      0 65001 65001 ?

*>                  0.0.0.0                  0         32768 ?
```

Lab 6-1: Establishing Overlapping VPNs

Complete this lab activity to practice what you learned in the related module.

Activity Objective

Your VPN customers want to exchange data between their central sites. You have decided to implement this request with an overlapping VPN topology.

In this activity, you will establish overlapping VPNs to support the needs of your customers. After completing this activity, you will be able to meet these objectives:

- Design a VPN solution
- Remove CE*x*1A and CE*x*2B from existing VRFs
- Configure new VRFs for CE*x*1A and CE*x*2B

Visual Objective

The figure illustrates what you will accomplish in this activity.

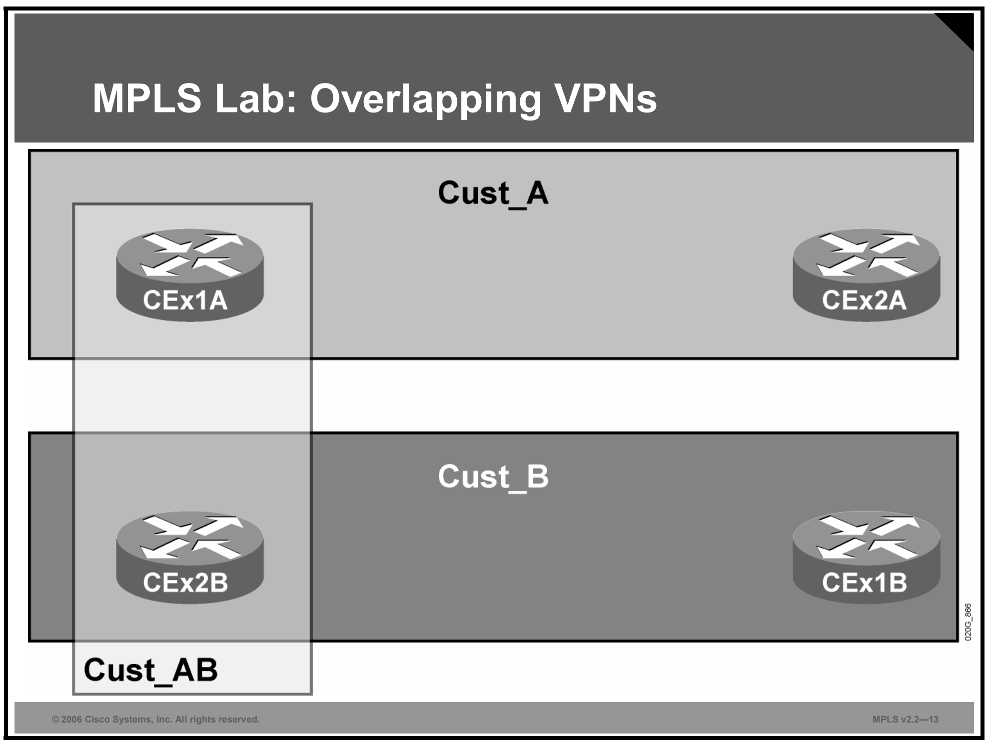

In this lab activity, you will establish overlapping VPNs with the following connectivity goals:

- Simple VPN communication
  - CE*x*1A and CE*x*2A can communicate.
  - CE*x*1B and CE*x*2B can communicate.
  - CE*x*1A and CE*x*1B cannot communicate.
  - CE*x*2A and CE*x*2B cannot communicate.
  - CE*x*1B and CE*x*2A cannot communicate.
- Overlapping VPN communication (CustAB)
  - CE*x*1A and CE*x*2B can communicate.

Required Resources

This is the resource that is required to complete this activity:

- Cisco IOS documentation

Command List

The commands that are used in this activity have been described in previous activities.

Task 1: Design Your VPN Solution

Site CE*x*1A cannot belong to the same VRF as the other *x*A sites. Similarly, site CE*x*2B cannot belong to the same VRF as the *x*B sites. Also, CE*x*1A and CE*x*2B cannot share the same VRF.

Activity Procedure

Complete these steps:

Step 1 On paper, allocate new RDs for VRFs to which CE*x*1A and CE*x*2B will be connected.

Step 2 A new RT is needed for the CustAB VPN. Coordinate the value of this RT with the other POP within your SP.

Note You could use *x*:11 as the RD for VRFs connected to CE*x*1A, and you could use *x*:21 as the RD for VRFs connected to CE*x*2B. You could use *x*:1001 as the RT for the CustAB VPN.

Activity Verification

You have completed this task when you attain this result:

- You have designed RDs and RTs for the new VRFs.

Task 2: Remove CE*x*1A and CE*x*2B from Existing VRFs

CE*x*1A and CE*x*2B must be migrated to new routing contexts. It is tempting to do this by merely changing the RDs and RTs of their existing VRF. However, this approach is not possible because the other VPN site, connected to the same PE router, is sharing those VRFs.

Note	When you enabled the backup link, you connected both CE*x*1A and CE*x*2A to PE*x*1. Therefore, if you change the routing context of customer A on PE*x*1, you will affect both CE*x*1A and CE*x*2A. This situation also holds true for CE*x*1B, CE*x*1B, and PE*x*2.

Sites CE*x*1A and CE*x*2B have to be migrated to new VRFs. All of the references to these sites must be removed from the existing routing protocol contexts.

In this task, you will remove the references to CE*x*1A and CE*x*2B.

Activity Procedure

Complete these steps:

Step 1 Remove the address family BGP neighbor relationship between CE*x*1A and CE*x*2B on their respective PE routers.

Step 2 Check any other references to CE*x*1A and CE*x*2B in their PE router configurations and, if required, remove them.

Activity Verification

You have completed this task when you attain these results:

- On the PE router, you have verified that the interface toward the CE router is no longer in the original VRF by using the **show ip vrf interfaces** command. This action should result in a printout similar to the one here:

```
PEx1#show ip vrf interfaces
Interface              IP-Address      VRF                      Protocol
Serial0/0.113          150.x.x1.50     CustA                    up
Serial0/0.102          150.x.x1.34     CustB                    up

PEx2#show ip vrf interfaces
Interface              IP-Address      VRF                      Protocol
Serial0/0.101          150.x.x2.18     CustA                    up
Serial0/0.113          150.x.x2.50     CustB                    up
```

- You have verified that the BGP neighbor relationship has been removed on the PE router with the **show ip bgp vpnv4 vrf summary** command. This action should give you a printout similar to the one here. Check the status of CE*x*1A and CE*x*2B in the printout.

```
PEx1#show ip bgp vpnv4 vrf CustA summary
BGP router identifier 192.168.x.17, local AS number 65001
BGP table version is 34, main routing table version 34
4 network entries using 847 bytes of memory
11 path entries using 704 bytes of memory
7 BGP path attribute entries using 1500 bytes of memory
1 BGP rrinfo entries using 24 bytes of memory
2 BGP AS-PATH entries using 48 bytes of memory
4 BGP extended community entries using 96 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 2139 total bytes of memory
BGP activity 51/29 prefixes, 69/43 paths, scan interval 15 secs
 Neighbor      V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
150.x.x1.49    4 650x1     976     979       34    0    0 00:29:12        4

PEx2#show ip bgp vpnv4 vrf CustB summary
BGP router identifier 192.168.x.33, local AS number 65001
BGP table version is 33, main routing table version 33
4 network entries using 605 bytes of memory
7 path entries using 448 bytes of memory
7 BGP path attribute entries using 1500 bytes of memory
1 BGP rrinfo entries using 24 bytes of memory
2 BGP AS-PATH entries using 48 bytes of memory
4 BGP extended community entries using 96 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 1642 total bytes of memory
BGP activity 122/102 prefixes, 160/138 paths, scan interval 15 secs

Neighbor     V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
150.x.x2.49  4 650x2    1477    1479       33    0    0 00:30:26        2
```

Task 3: Configure New VRFs for CE*x*1A and CE*x*2B

In this task, you will create the new VRFs for CE*x*1A and CE*x*2B.

Activity Procedure

Complete these steps:

Step 1 Create the new VRFs for CE*x*1A and CE*x*2B on their PE router with the **ip vrf** command.

Step 2 Assign new RDs to the newly created VRFs with the **rd** command.

Step 3 Assign proper import and export RTs to the newly created VRFs with the **route-target** command.

Step 4 Reestablish BGP routing between the PE routers and the CE routers. Please refer to Lab 5-4, "Running BGP Between PE and CE Routers," if you need more details.

Activity Verification

You have completed this task when you attain these results:

- On the PE router, you have verified that the interface toward the CE router is in the proper VRF by using the **show ip vrf interfaces** command. This action should result in a printout similar to the one here:

```
PEx1#show ip vrf interfaces
Interface              IP-Address      VRF              Protocol
Serial0/0.113          150.x.x1.50     CustA         up
Serial0/0.101          150.x.x1.18     CustAB        up
Serial0/0.102          150.x.x1.34     CustB         up

PEx2#show ip vrf interfaces
Interface              IP-Address      VRF              Protocol
Serial0/0.101          150.x.x2.18     CustA         up
Serial0/0.102          150.x.x2.34     CustAB        up
Serial0/0.113          150.x.x2.50     CustB         up
```

- You have verified the BGP neighbors on the PE router with the **show ip bgp vpnv4 vrf summary** command. This action should result in a printout similar to the one here. Check the status of CE*x*1A and CE*x*2B in the printout.

```
PEx1#show ip bgp vpnv4 vrf CustAB summary
BGP router identifier 192.168.x.17, local AS number 65001
BGP table version is 49, main routing table version 49
10 network entries using 1210 bytes of memory
10 path entries using 640 bytes of memory
8 BGP path attribute entries using 480 bytes of memory
2 BGP AS-PATH entries using 48 bytes of memory
4 BGP extended community entries using 96 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
```

```
BGP using 2474 total bytes of memory
BGP activity 57/35 prefixes, 75/49 paths, scan interval 15 secs

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
150.x.x1.17  4 650x1      53      54       49    0    0 00:48:43        3

PEx2#show ip bgp vpnv4 vrf CustAB summary
BGP router identifier 192.168.x.33, local AS number 65001
BGP table version is 56, main routing table version 56
10 network entries using 1210 bytes of memory
10 path entries using 640 bytes of memory
8 BGP path attribute entries using 480 bytes of memory
2 BGP AS-PATH entries using 48 bytes of memory
4 BGP extended community entries using 96 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 2068 total bytes of memory
BGP activity 130/110 prefixes, 168/146 paths, scan interval 15 secs

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
150.x.x2.33  4 650x2       9      10       56    0    0 00:04:17        3
```

- You have checked the BGP routing table in the new VRF with the **show ip bgp vpnv4 vrf** command. You should see routes from CE*x*1A or CE*x*2B and routes imported from other VRFs. Use the AS path to work out which routes belong to which CE router. Routes announced by CE*x*1A should have 650*x*1 in the AS path, and routes announced by CE*x*2B should have 650*x*2 in the AS path.

```
PEx1#show ip bgp vpnv4 vrf CustAB
BGP table version is 49, local router ID is 192.168.x.17
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: x:11 (default for vrf CustAB)
*> 10.1.x1.16/28     150.x.x1.17              0             0 650x1 ?
*> 10.1.x1.49/32     150.x.x1.17              0             0 650x1 ?
*>i10.1.x2.16/28     192.168.x.33             0    100      0 650x1 ?
*>i10.1.x2.49/32     192.168.x.33             0    100      0 650x2 ?
*>i10.2.x2.16/28     192.168.x.33             0    100      0 650x2 ?
*>i10.2.x2.49/32     192.168.x.33             0    100      0 650x1 ?
r> 150.x.x1.16/28   150.x.x1.17              0             0 650x1 ?
*>i150.x.x1.48/28   192.168.x.33             0    100      0 650x1 ?
*>i150.x.x2.16/28   192.168.x.33             0    100      0 650x1 ?
*>i150.x.x2.32/28   192.168.x.33             0    100      0 650x2 ?
```

```
PEx2#show ip bgp vpnv4 vrf CustAB
BGP table version is 95, local router ID is 192.168.x.33
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: x:21 (default for vrf CustAB)
*>i10.1.x1.16/28    192.168.x.17             0    100      0 650x1 ?
*>i10.1.x1.49/32    192.168.x.17             0    100      0 650x1 ?
*>i10.2.x1.16/28    192.168.x.17             0    100      0 650x2 ?
*>i10.2.x1.49/32    192.168.x.17             0    100      0 650x2 ?
*> 10.2.x2.16/28    150.x.x2.33              0             0 650x2 ?
*> 10.2.x2.49/32    150.x.x2.33              0             0 650x2 ?
*>i150.x.x1.16/28   192.168.x.17             0    100      0 650x1 ?
*>i150.x.x1.32/28   192.168.x.17             0    100      0 650x2 ?
r> 150.x.x2.32/28   150.x.x2.33              0             0 650x2 ?
*>i150.x.x2.48/28   192.168.x.17             0    100      0 650x2 ?
```

- You have connected to CE*x*1A and performed **ping** and **trace** tests to the loopback address of CE*x*2B (or vice versa). The other router should be reachable. For subgroup B, perform the test in the other direction.

```
CEx1A#ping 10.2.x2.49

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.2.x2.49, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 148/148/149 ms

CEx1A#trace 10.2.x2.49
Type escape sequence to abort.
Tracing the route to 10.2.x2.49
  1 150.x.x1.18 16 msec 16 msec 12 msec
  2 150.x.x2.34 [AS 650x2] [MPLS: Label 17 Exp 0] 116 msec 116 msec 116 msec
  3 150.x.x2.33 [AS 650x2] 72 msec 77 msec *
```

- From CE*x*1A, perform a **ping** test to the loopback address of CE*x*1B (or vice versa). The other router should not be reachable.

```
CEx1A#ping 10.2.x1.49

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.2.x1.49, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

- Connect to CE*x*2A and try to ping CE*x*2B or CE*x*1B. Those routers should not be reachable from CE*x*2A.

```
CEx2A#ping 10.2.x2.49
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.2.x2.49, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

CEx2A#ping 10.2.x1.49
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.2.x1.49, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Lab 6-2: Merging Service Providers

Complete this lab activity to practice what you learned in the related module.

Activity Objective

Your small service provider is merging with several other small service providers. To accomplish this consolidation, a new central P router (P1) has been installed and configured. Frame Relay connectivity has been provided from each local P*x*1 and P*x*2 router to P1. In addition, the core IGP is being converted from EIGRP to IS-IS.

In this activity, you will merge your small service provider with several other small service providers. After completing this activity, you will be able to meet these objectives:

- Establish connectivity with the central P1 router
- Convert the core IGP from EIGRP to IS-IS
- Enable MPLS LDP connectivity with the central P router
- Enable IBGP connectivity between all PE routers

Visual Objective

You will configure your PE router, and the directly connected P router. P1 has been preconfigured.

The figure illustrates what you will accomplish in this activity.

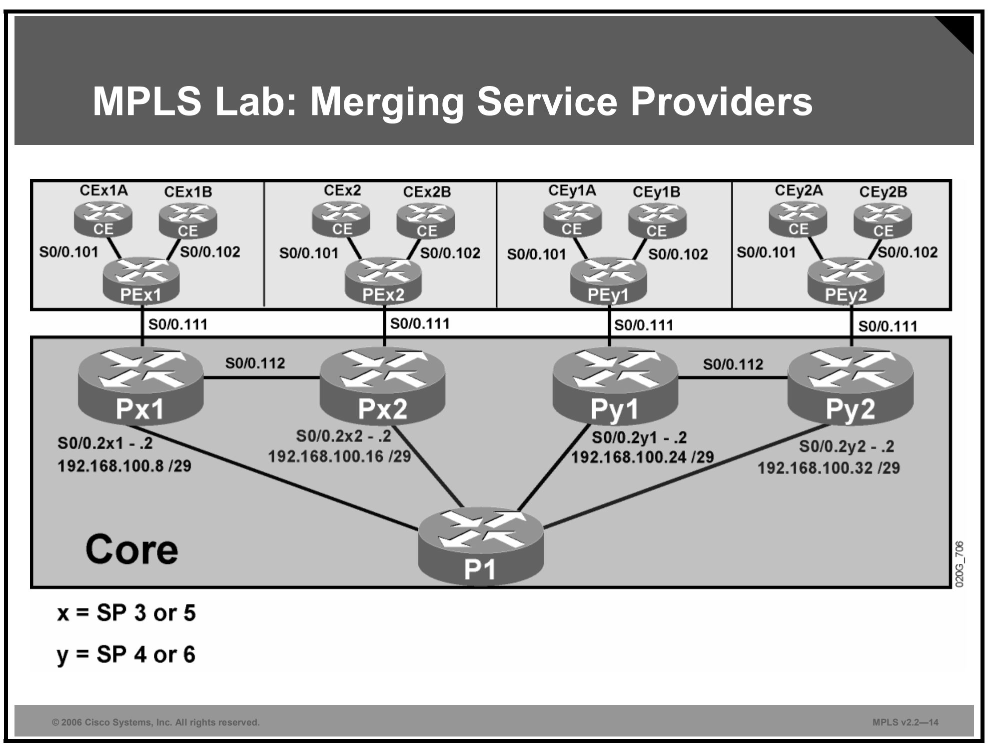

Required Resources

This is the resource that is required to complete this activity:

- Cisco IOS documentation

Command List

The table describes the commands that are used in this activity.

Commands for Merging Service Providers

Command	Description
router isis area-tag	To enable the IS-IS routing protocol and to specify an IS-IS process, use the **router isis** command in global configuration mode. To disable IS-IS routing, use the **no** form of this command.
net network-entity-title	To configure an IS-IS network entity title (NET) for a Connectionless Network Service (CLNS) routing process, use the **net** command in router configuration mode. To remove a NET, use the **no** form of this command.
isis circuit-type {**level-1** | **level-1-2** | **level-2-only**}	To configure the type of adjacency, use the **isis circuit-type** interface configuration command. To reset the circuit type to Level l and Level 2, use the **no** form of this command.
metric-style wide [**transition**] [**level-1** | **level-2** | **level-1-2**]	To configure a router running IS-IS so that it generates and accepts only new-style type, TLV objects, use the **metric-style wide** command in router configuration mode. To disable this function, use the **no** form of this command.

Task 1: Enable Connectivity with the Central P Router

In this task, you will enable the Frame Relay link between your P routers and P1, and then enable LDP connectivity between the two routers.

Activity Procedure

Complete these steps:

Step 1 Configure IP addresses and DLCIs on your P router on the interface to P1 using the parameters in the table.

Note The parameters are configured on the P routers of the SP and not the PE routers.

IP Address and DLCI Configuration Parameters

Router	Subinterface	DLCI	IP Address
P31	S0/0.231	231	192.168.100.10/29
P32	S0/0.232	232	192.168.100.18/29
P41	S0/0.241	241	192.168.100.26/29
P42	S0/0.242	242	192.168.100.34/29

Router	Subinterface	DLCI	IP Address
P51	S0/0.251	251	192.168.100.42/29
P52	S0/0.252	252	192.168.100.50/29
P61	S0/0.261	261	192.168.100.58/29
P62	S0/0.262	262	192.168.100.66/29

Activity Verification

You have completed this task when you attain this result:

- On your P router, you have used the **show interface** command to verify that the new interfaces are operational.

Task 2: Migrate the Core to IS-IS

Because a link-state protocol is more scalable than a distance vector protocol, the service provider has decided to migrate the core to IS-IS. The P1 router has already been migrated. Your POP is responsible for the migration of all of your assigned routers. POP 1 will migrate PE*x*1 and P*x*1. POP 2 will migrate PE*x*2 and P*x*2.

Activity Procedure

Complete these steps:

Step 1 Disable EIGRP as the core IGP on your assigned routers.

Step 2 Enable IS-IS as the core IGP using the parameters detailed in the table.

IS-IS Parameters

Router ID	NET	Remarks
PE*x*1	net 49.0001.0000.0000.01*x*1.00	Where *x* = the SP number
PE*x*2	net 49.0001.0000.0000.01*x*2.00	
P*x*1	net 49.0001.0000.0000.02*x*1.00	
P*x*2	net 49.0001.0000.0000.02*x*2.00	

Note Ensure that the **metric-style** command is set to **wide**, the **is-type** command is set to **level-2-only**, and IS-IS has been enabled on the active serial interfaces that are supporting the core MPLS.

Activity Verification

You have completed this task when you attain these results:

- You have used the **show ip protocols** command to verify that IS-IS is active and enabled on all appropriate interfaces on the PE routers.

```
PEx1#show ip protocols
Routing Protocol is "bgp 65001"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  IGP synchronization is disabled
  Automatic route summarization is disabled
  Neighbor(s):
    Address          FiltIn FiltOut DistIn DistOut Weight RouteMap
    192.168.x.33
  Maximum path: 1
  Routing Information Sources:
    Gateway         Distance      Last Update
  Distance: external 20 internal 200 local 200
Routing Protocol is "isis"
  Invalid after 0 seconds, hold down 0, flushed after 0
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Redistributing: isis
  Address Summarization:
    None
  Maximum path: 4
  Routing for Networks:
    Serial0/0.111
    Loopback0
  Routing Information Sources:
    Gateway         Distance      Last Update
    192.168.x.97         115      00:08:32
    192.168.y.97         115      00:08:32
    192.168.x.81         115      00:08:22
    192.168.y.81         115      00:08:32
    192.168.x.33         115      00:08:22
    192.168.y.33         115      00:08:32
    192.168.y.17         115      00:08:32
    192.168.100.129      115      00:08:32
  Distance: (default is 115)
```

Note The IS-IS gateways should show the loopback addresses of the PE and P routers. In these results, *x* is your SP number, and *y* is the number of the other SP.

Depending on how far along the other SP is in configuring its devices, you may not see the *y* routes.

- You have used the **show ip route** command on the PE routers to verify that all routers are sending and receiving the appropriate prefixes.

```
PEx1#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, 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

     192.168.y.0/24 is variably subnetted, 7 subnets, 2 masks
i L2    192.168.y.97/32 [115/40] via 192.168.x.50, Serial0/0.111
i L2    192.168.y.112/28 [115/40] via 192.168.x.50, Serial0/0.111
i L2    192.168.y.64/28 [115/40] via 192.168.x.50, Serial0/0.111
i L2    192.168.y.81/32 [115/40] via 192.168.x.50, Serial0/0.111
i L2    192.168.y.33/32 [115/50] via 192.168.x.50, Serial0/0.111
i L2    192.168.y.48/28 [115/40] via 192.168.x.50, Serial0/0.111
i L2    192.168.y.17/32 [115/50] via 192.168.x.50, Serial0/0.111
     192.168.100.0/24 is variably subnetted, 5 subnets, 2 masks
i L2    192.168.100.8/29 [115/20] via 192.168.x.50, Serial0/0.111
i L2    192.168.100.24/29 [115/30] via 192.168.x.50, Serial0/0.111
i L2    192.168.100.16/29 [115/30] via 192.168.x.50, Serial0/0.111
i L2    192.168.100.32/29 [115/30] via 192.168.x.50, Serial0/0.111
i L2    192.168.100.129/32 [115/30] via 192.168.x.50, Serial0/0.111
     192.168.x.0/24 is variably subnetted, 7 subnets, 2 masks
i L2    192.168.x.97/32 [115/30] via 192.168.x.50, Serial0/0.111
i L2    192.168.x.112/28 [115/20] via 192.168.x.50, Serial0/0.111
i L2    192.168.x.64/28 [115/30] via 192.168.x.50, Serial0/0.111
i L2    192.168.x.81/32 [115/20] via 192.168.x.50, Serial0/0.111
i L2    192.168.x.33/32 [115/40] via 192.168.x.50, Serial0/0.111
C       192.168.x.48/28 is directly connected, Serial0/0.111
C       192.168.x.17/32 is directly connected, Loopback0
```

Note	In these results, *x* is your SP number, and *y* is the number of the other SP. Depending on how far along the other SP is in configuring its devices, you may not see the *y* routes.

■ You have used the **show ip protocols** command to verify that IS-IS is active and enabled on all appropriate interfaces on the P routers.

```
Px1#show ip protocols
Routing Protocol is "isis"
  Invalid after 0 seconds, hold down 0, flushed after 0
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Redistributing: isis
  Address Summarization:
    None
  Maximum path: 4
  Routing for Networks:
    Serial0/0.111
    Serial0/0.112
    Serial0/0.2x1
    Loopback0
  Routing Information Sources:
    Gateway         Distance      Last Update
    192.168.x.97         115      00:00:59
    192.168.y.97         115      00:12:07
    192.168.y.81         115      00:14:36
    192.168.x.33         115      00:10:04
    192.168.y.33         115      00:12:07
    192.168.x.17         115      00:00:59
    192.168.y.17         115      00:12:17
    192.168.100.129      115      00:00:59
  Distance: (default is 115)
```

Note In these results, *x* is your SP number, and *y* is the number of the other SP.

■ You have used the **show ip route** command on the P routers to verify that all routers are sending and receiving the appropriate prefixes.

```
Px1#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, 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

     192.168.y.0/24 is variably subnetted, 7 subnets, 2 masks
i L2    192.168.y.97/32 [115/30] via 192.168.100.9, Serial0/0.231
i L2    192.168.y.112/28 [115/30] via 192.168.100.9, Serial0/0.231
```

```
i L2    192.168.y.64/28 [115/30] via 192.168.100.9, Serial0/0.231
i L2    192.168.y.81/32 [115/30] via 192.168.100.9, Serial0/0.231
i L2    192.168.y.33/32 [115/40] via 192.168.100.9, Serial0/0.231
i L2    192.168.y.48/28 [115/30] via 192.168.100.9, Serial0/0.231
i L2    192.168.y.17/32 [115/40] via 192.168.100.9, Serial0/0.231
     192.168.100.0/24 is variably subnetted, 5 subnets, 2 masks
C       192.168.100.8/29 is directly connected, Serial0/0.231
i L2    192.168.100.24/29 [115/20] via 192.168.100.9, Serial0/0.2x1
i L2    192.168.100.16/29 [115/20] via 192.168.100.9, Serial0/0.2x1
                          [115/20] via 192.168.x.114, Serial0/0.112
i L2    192.168.100.32/29 [115/20] via 192.168.100.9, Serial0/0.2x1
i L2    192.168.100.129/32 [115/20] via 192.168.100.9, Serial0/0.2x1
     192.168.x.0/24 is variably subnetted, 7 subnets, 2 masks
i L2    192.168.x.97/32 [115/20] via 192.168.x.114, Serial0/0.112
C       192.168.x.112/28 is directly connected, Serial0/0.112
i L2    192.168.x.64/28 [115/20] via 192.168.x.114, Serial0/0.112
C       192.168.x.81/32 is directly connected, Loopback0
i L2    192.168.x.33/32 [115/30] via 192.168.x.114, Serial0/0.112
C       192.168.x.48/28 is directly connected, Serial0/0.111
i L2    192.168.x.17/32 [115/20] via 192.168.x.49, Serial0/0.111
```

Note In these results, *x* is your SP number, and *y* is the number of the other SP.

Task 3: Enable MPLS LDP Connectivity with the Central P Router

In this task you will enable LDP connectivity between your routers and P1.

Activity Procedure

Complete this step:

Step 1 Enable LDP on the subinterface that you have created on P*x*1 or P*x*2 router.

Activity Verification

You have completed this task when you attain these results:

- On your P router, you have verified that an LDP neighbor relationship has been established between your P router and P1.

```
Px1#show mpls ldp neighbor
    Peer LDP Ident: 192.168.x.17:0; Local LDP Ident 192.168.x.81:0
        TCP connection: 192.168.x.17.646 - 192.168.x.81.44875
        State: Oper; Msgs sent/rcvd: 135/130; Downstream
        Up time: 01:36:03
        LDP discovery sources:
          Serial0/0.111, Src IP addr: 192.168.x.49
        Addresses bound to peer LDP Ident:
```

```
      192.168.x.17    192.168.x.49
  Peer LDP Ident: 192.168.x.97:0; Local LDP Ident 192.168.x.81:0
    TCP connection: 192.168.x.97.54451 - 192.168.x.81.646
    State: Oper; Msgs sent/rcvd: 136/136; Downstream
    Up time: 01:36:03
    LDP discovery sources:
      Serial0/0.112, Src IP addr: 192.168.x.114
    Addresses bound to peer LDP Ident:
      192.168.x.97    192.168.x.66    192.168.x.114   192.168.100.34
  Peer LDP Ident: 192.168.100.129:0; Local LDP Ident 192.168.x.81:0
    TCP connection: 192.168.100.129.36547 - 192.168.x.81.646
    State: Oper; Msgs sent/rcvd: 24/30; Downstream
    Up time: 00:02:25
    LDP discovery sources:
      Serial0/0.2x1, Src IP addr: 192.168.100.25
    Addresses bound to peer LDP Ident:
      192.168.100.129 201.202.20.1    201.202.21.1    201.202.22.1
      201.202.23.1    201.202.24.1    201.202.25.1    192.168.100.9
      192.168.100.17  192.168.100.25  192.168.100.33
```

- On your PE router, you have verified that labels are being received from the other POPs.

```
PEx1#show mpls forwarding-table
Local  Outgoing    Prefix            Bytes tag  Outgoing   Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
18     Pop tag     192.168.100.8/29  0          Se0/0.111  point2point
19     Pop tag     192.168.x.112/28  0          Se0/0.111  point2point
20     Pop tag     192.168.x.81/32   0          Se0/0.111  point2point
21     18          192.168.100.16/29 0          Se0/0.111  point2point
22     20          192.168.x.97/32   0          Se0/0.111  point2point
25     19          192.168.x.64/28   0          Se0/0.111  point2point
26     17          192.168.100.24/29 0          Se0/0.111  point2point
27     22          192.168.100.32/29 0          Se0/0.111  point2point
28     23          192.168.100.129/32   \
                                     0          Se0/0.111  point2point
29     Untagged    10.1.x1.16/28[V]  0          Se0/0.101  point2point
30     Untagged    10.1.x1.49/32[V]  0          Se0/0.101  point2point
31     Aggregate   150.x.x1.16/28[V] 520
32     Untagged    10.2.x1.16/28[V]  0          Se0/0.102  point2point
33     16          192.168.y.112/28  0          Se0/0.111  point2point
34     21          192.168.y.81/32   0          Se0/0.111  point2point
35     24          192.168.y.48/28   0          Se0/0.111  point2point
36     Untagged    10.2.x1.49/32[V]  0          Se0/0.102  point2point
37     Aggregate   150.x.x1.32/28[V] 0
38     Untagged    150.x.x2.48/28[V] 0          Se0/0.102  point2point
39     25          192.168.y.97/32   0          Se0/0.111  point2point
40     26          192.168.y.64/28   0          Se0/0.111  point2point
41     27          192.168.y.33/32   0          Se0/0.111  point2point
```


42	28	192.168.y.17/32	0	Se0/0.111	point2point
43	30	192.168.x.33/32	0	Se0/0.111	point2point

Note In these results, *x* is your SP number, and *y* is the number of the other SP. The "[V]" indicates a VPN prefix.

Task 4: Enable IBGP Connectivity for All PE Routers

At this point, you have established LDP connectivity for all of the PE and P routers in your new service provider environment, but you have not yet established BGP connectivity. You now need to establish IBGP connectivity for your PE routers.

There are two methods that you can implement. The first is to use the **bgp neighbor** command to add a neighbor relationship between each of the routers, but this approach would entail a substantial configuration effort.

The second method is to implement route reflectors. To this end, P1 has been configured as a BGP route reflector. However, to take advantage of this fact, you will need to remove the neighbor relationship between your two PE routers and make them neighbors of P1.

Note The loopback address for P1 is 192.168.100.129 with AS 65001. Ensure that your update source is also your loopback interface.

POP 1 will configure PE*x*1, and POP 2 will configure PE*x*2.

Activity Procedure

Complete these steps:

Step 1 Remove the neighbor relationship between your PE router and the PE router in your remote POP.

Step 2 Configure your PE router as a neighbor of P1.

What routes do you expect to see in VRF CustA?

____

What routes do you expect to see in VRF CustB?

____

What routes do you expect to see in VRF CustAB?

____

Activity Verification

You have completed this task when you attain these results:

- On your PE routers, you have checked BGP connectivity to all POPs with the **show ip bgp summary** and **show ip bgp neighbor** commands on CE routers.

```
PEx1#show ip bgp summary
Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down State/PfxRcd
192.168.100.129 4 65001      18      16        4    0    0 00:04:26        1
```

- You have verified the per-VRF BGP table for your customer on your PE routers with the **show ip bgp vpnv4 vrf** command. You should still see that the BGP routes coming from the CE routers are being selected as the best routes for those destinations.

```
PEx1#show ip bgp vpnv4 vrf CustA
BGP table version is 198, local router ID is 192.168.x.17
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: 1:10 (default for vrf CustA)
*> 10.1.x1.16/28    150.x.x1.17              0             0 650x1 ?
*> 10.1.x1.49/32    150.x.x1.17              0             0 650x1 ?
*>i10.1.x2.16/28    192.168.x.33             0    100      0 650x1 ?
*                   150.x.x1.49            200             0 650x1 ?
*>i10.1.x2.49/32    192.168.x.33             0    100      0 650x1 ?
*                   150.x.x1.49            200             0 650x1 ?
*> 150.x.x1.16/28   150.x.x1.17              0             0 650x1 ?
r>i150.x.x1.48/28   192.168.x.33             0    100      0 650x1 ?
r                   150.x.x1.49            200             0 650x1 ?
*>i150.x.x2.16/28   192.168.x.33             0    100      0 650x1 ?
*                   150.x.x1.49            200             0 650x1 ?
```

```
PEx2#sh ip bgp vp vrf CustB
BGP table version is 38, local router ID is 192.168.x.33
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: x:20 (default for vrf CustB)
*> 10.2.x2.16/28    150.x.x2.33              0             0 650x2 ?
*> 10.2.x2.49/32    150.x.x2.33              0             0 650x2 ?
*>i10.2.x1.16/28    192.168.x.17             0    100      0 650x2 ?
*                   150.x.x2.49            200             0 650x2 ?
*>i10.2.x1.49/32    192.168.x.17             0    100      0 650x2 ?
*                   150.x.x2.49            200             0 650x2 ?
```



```
> 150.x.x2.32/28    150.x.x2.33              0             0 650x2 ?
r>i150.x.x2.48/28   192.168.x.17             0    100      0 650x2 ?
r                   150.x.x2.49            200             0 650x2 ?
*>i150.x.x1.32/28   192.168.x.17             0    100      0 650x2 ?
*                   150.x.x2.49            200             0 650x2 ?
PE52>
```

- You have verified the per-VRF table for your customer on your PE routers with the **show ip route vrf** command. You should still see only the routes coming from the CE routers being selected.

```
PEx1#show ip route vrf CustA

Routing Table: CustA
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, 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, 4 subnets, 2 masks
B       10.1.x1.16/28 [20/0] via 150.x.x1.17 (CustAB), 04:51:01
B       10.1.x2.16/28 [200/0] via 192.168.x.33, 00:27:42
B       10.1.x1.49/32 [20/0] via 150.x.x1.17 (CustAB), 04:51:01
B       10.1.x2.49/32 [200/0] via 192.168.x.33, 00:27:42
     150.x.0.0/28 is subnetted, 3 subnets
B       150.x.x2.16 [200/0] via 192.168.x.33, 00:27:42
B       150.x.x1.16 [20/0] via 150.x.x1.17 (CustAB), 04:51:01
C       150.x.x1.48 is directly connected, Serial0/0.113

PEx1#show ip route vrf CustB

Routing Table: CustB
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, 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, 4 subnets, 2 masks
```

```
B       10.2.x1.16/28 [20/0] via 150.x.x1.33, 05:53:38
B       10.2.x2.16/28 [200/0] via 192.168.x.33, 00:30:58
B       10.2.x1.49/32 [20/0] via 150.x.x1.33, 05:53:11
B       10.2.x2.49/32 [200/0] via 192.168.x.33, 00:30:58
     150.x.0.0/28 is subnetted, 3 subnets
B       150.x.x2.48 [20/0] via 150.x.x1.33, 05:53:11
B       150.x.x2.32 [200/0] via 192.168.x.33, 00:30:58
C       150.x.x1.32 is directly connected, Serial0/0.102

PEx1#show ip route vrf CustAB

Routing Table: CustAB
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, 6 subnets, 2 masks
B       10.1.x1.16/28 [20/0] via 150.x.x1.17, 03:00:00
B       10.2.x2.49/32 [200/0] via 192.168.x.33, 00:13:46
B       10.1.x2.49/32 [200/0] via 192.168.x.33, 00:13:46
B       10.1.x1.49/32 [20/0] via 150.x.x1.17, 03:00:00
B       10.2.x2.16/28 [200/0] via 192.168.x.33, 00:13:46
B       10.1.x2.16/28 [200/0] via 192.168.x.33, 00:13:46
     150.x.0.0/28 is subnetted, 4 subnets
B       150.x.x2.32 [200/0] via 192.168.x.33, 00:13:47
C       150.x.x1.16 is directly connected, Serial0/0.101
B       150.x.x2.16 [200/0] via 192.168.x.33, 00:13:47
B       150.x.x1.48 [200/0] via 192.168.x.33, 00:13:48
PEx1#
```

Lab 6-3: Establishing a Common Services VPN

The new MPLS VPN infrastructure can be used to implement a new approach to managed CE router services, where the central NMS can monitor all CE routers through a dedicated VPN.

The NMS VPN should provide connectivity only between the NMS and a single IP address on the CE router that is used for network management purposes.

In this activity, your SP has established a network management center using a VPN between the loopback interfaces of the CE routers and the NMS router. You will establish connectivity only between the NMS and the CE router loopback interfaces with a /32 subnet mask.

Complete this lab activity to practice what you learned in the related module.

Activity Objective

In this activity, you will establish a network management VPN between the loopback interfaces of the CE routers and the NMS router. After completing this activity, you will be able to meet these objectives:

- Implement a network management VPN between the management VRF and customer VRFs by configuring proper RTs

Visual Objective

The figure illustrates what you will accomplish in this activity.

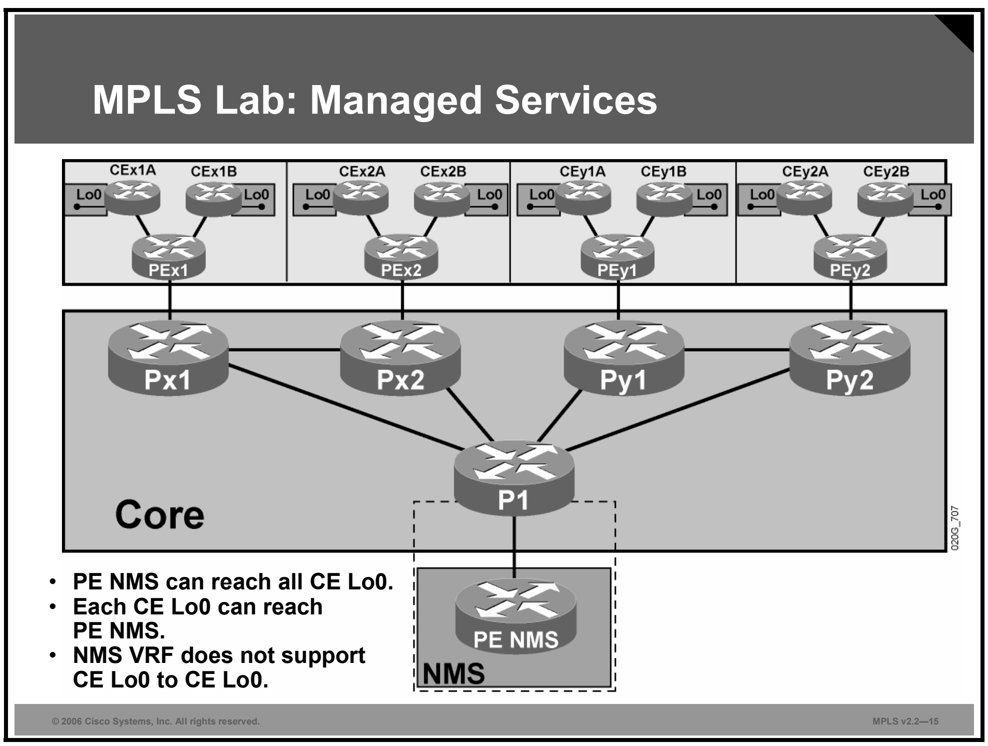

Note The NMS router is simulated by a loopback on the P1 router.

Required Resources

This is the resource that is required to complete this activity:

- Cisco IOS documentation

Command List

The table describes the commands that are used in this activity.

Network Management VPN Commands

Command	Description
export map *name*	Specifies a VRF export route map
route-map *name* **permit** *seq*	Creates a route map entry
set extcommunity rt *value* **additive**	Appends the specified RT to a route matched with the **match** command

Task 1: Establish Connectivity Between the NMS VRF and Other VRFs

The network management VPN is a common services VPN. Therefore, two RTs are needed for the VPN: the server RT and the client RT. On the PE router supporting the NMS, a VRF for the network management VPN and associated RD are also needed. Here are the relevant parts of the configuration on the NMS PE router:

Note This configuration resides on the P1 router, which in this exercise is simulating the central service PE router.

```
! Create the NMS VRF
!
ip vrf NMS
 rd 101:500
 route-target export 101:500
 route-target import 101:500
 route-target import 101:501
```

Note You will need to configure the VRF of your customer only on the local PE router to match the RT used by the NMS VPN.

To establish connectivity between the NMS VRF and the customer VRF, you must attach the client RT to the CE router loopback addresses when the routes are exported from the customer VRF. You also need to import routes from the NMS router into all customer VRFs.

Activity Procedure

Complete these steps:

Step 1 Create an IP access list that will match the CE router loopback addresses.

Step 2 Create a route map that will match the CE router loopback addresses with the access list and append the client RT to those routes.

Step 3 Apply the route map to routes exported from the customer VRF with the **export map** command.

Step 4 Import NMS routes into the customer VRF by specifying the proper import RT.

What routes do you expect to see in VRF CustA on your PE?

What routes do you expect to see in VRF NMS on your PE?

What routes do you expect to see in VRF NMS on P1?

Activity Verification

You have completed this task when you attain these results:

- You have verified that the proper RTs are appended to the routes toward the CE router loopback addresses by using the **show ip bgp vpnv4 vrf** *name prefix* command. This action should result in a printout similar to this example:

```
PEx1#show ip bgp vpnv4 vrf CustA 10.1.x1.49
BGP routing table entry for 1:10:10.1.x1.49/32, version 46
Paths: (1 available, best #1, table CustA)
  Advertised to non peer-group peers:
  150.x.x1.49
  650x1, imported path from x:11:10.1.x1.49/32
    150.x.x1.17 from 150.x.x1.17 (10.1.x1.49)
      Origin incomplete, metric 0, localpref 100, valid, external, best
      Extended Community: RT:x:10 RT:x:1001 RT:101:501
```

Note If after a few minutes you do not see RT:101:501 in the extended community list on the PE router, you may need to use the **clear ip bgp *** command to reset the BGP session.

- You have verified that the proper routes are in the VRF CustA on your PE router. You should now see the simulated NMS address and the prefixes for your CustA routes.

```
PEx1#show ip bgp vp vrf CustA
BGP table version is 74, local router ID is 192.168.3.17
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: 3:10 (default for vrf CustA)
*> 10.1.x1.16/28    150.x.x1.17              0             0 65031 ?
*> 10.1.x1.49/32    150.x.x1.17              0             0 65031 ?
*>i10.1.x2.16/28    192.168.x.33             0    100      0 65031 ?
*                   150.x.x1.49            200             0 65031 ?
*>i10.1.x2.49/32    192.168.x.33             0    100      0 65031 ?
*                   150.x.x1.49            200             0 65031 ?
*>i10.10.10.49/32   192.168.100.129          0    100      0 ?
*> 150.x.x1.16/28   150.x.x1.17              0             0 65031 ?
r>i150.x.x1.48/28   192.168.x.33             0    100      0 65031 ?
r                   150.x.x1.49            200             0 65031 ?
*>i150.x.x2.16/28   192.168.x.33             0    100      0 65031 ?
*                   150.x.x1.49            200             0 65031 ?
PEx1#
```

- You have verified that no routes are in the VRF NMS on your PE router. You should not see any routes because you do not have interfaces in that VRF, and the VRF is not configured on your router.

```
PEx1#show ip bgp vp vrf NMS
PEx1#

PEx1#show ip vrf
  Name                             Default RD          Interfaces
  CustA                            x:10                Se0/0.113
  CustAB                           x:11                Se0/0.101
  CustB                            x:20                Se0/0.102
```

- Using an **extended ping** command, you have verified that you can ping from the loopback address of the managed CE router to the loopback address of the NMS CE router (10.10.10.49).

```
CEx1A#ping
Protocol [ip]:
Target IP address: 10.10.10.49
Repeat count [5]:
Datagram size [100]:
Timeout in seconds [2]:
Extended commands [n]: y
Source address or interface: loopback0
```

```
Type of service [0]:
Set DF bit in IP header? [no]:
Validate reply data? [no]:
Data pattern [0xABCD]:
Loose, Strict, Record, Timestamp, Verbose[none]:
Sweep range of sizes [n]:
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.10.49, timeout is 2 seconds:
Packet sent with a source address of 10.1.x1.49
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 88/114/212 ms
```

- Using an **extended ping** command, you have verified that you cannot ping from the Ethernet address of the managed CE router to the loopback address of the NMS CE router (10.10.10.49).
- You have verified that your CE router is seeing only prefixes within your VPN and that no prefixes are being leaked from other VPNs.

```
CEx1A>sh ip bgp
BGP table version is 12, local router ID is 10.1.x1.49
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
*> 10.1.x1.16/28    0.0.0.0                  0         32768 ?
*> 10.1.x1.49/32    0.0.0.0                  0         32768 ?
*> 10.1.x2.16/28    150.x.x1.18                            0 65001 65001 ?
*> 10.1.x2.49/32    150.x.x1.18                            0 65001 65001 ?
*> 10.2.x2.16/28    150.x.x1.18                            0 65001 65032 ?
*> 10.2.x2.49/32    150.x.x1.18                            0 65001 65032 ?
*> 10.10.10.49/32   150.x.x1.18                            0 65001 ?
*> 150.x.x1.16/28   0.0.0.0                  0         32768 ?
*> 150.x.x1.48/28   150.x.x1.18                            0 65001 65001 ?
*> 150.x.x2.16/28   150.x.x1.18                            0 65001 65001 ?
*> 150.x.x2.32/28   150.x.x1.18                            0 65001 65032 ?

CEx1A>sh ip ro
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, 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, 7 subnets, 2 masks
C       10.1.x1.16/28 is directly connected, Ethernet0/0
B       10.2.x2.49/32 [20/0] via 150. x.x1.18, 00:05:17
B       10.1.x2.49/32 [20/0] via 150. x.x1.18, 00:05:17
C       10.1.x1.49/32 is directly connected, Loopback0
B       10.2.x2.16/28 [20/0] via 150. x.x1.18, 00:05:17
B       10.10.10.49/32 [20/0] via 150.x.x1.18, 00:05:17
B       10.1.x2.16/28 [20/0] via 150. x.x1.18, 00:05:17
     150.x.0.0/28 is subnetted, 4 subnets
B       150.x.x2.32 [20/0] via 150.3.31.18, 00:05:17
C       150.x.x1.16 is directly connected, Serial0/0.101
B       150.x.x2.16 [20/0] via 150.3.31.18, 00:05:17
B       150.x.x1.48 [20/0] via 150.3.31.18, 00:05:17
CE31A>
```

- You have verified that the P router has only the management prefixes in the NMS VRF.

```
P1>sh ip bgp vpn vrf NMS
BGP table version is 135, local router ID is 201.202.25.1
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: 101:500 (default for vrf NMS)
*>i10.1.x1.49/32    192.168.x.17             0    100      0 65031 ?
*>i10.1.x2.49/32    192.168.x.33             0    100      0 65031 ?
*>i10.1.y1.49/32    192.168.y.17             0    100      0 65041 ?
*>i10.1.y2.49/32    192.168.y.33             0    100      0 65041 ?
*>i10.2.x1.49/32    192.168.x.17             0    100      0 65032 ?
*>i10.2.x2.49/32    192.168.x.33             0    100      0 65032 ?
*>i10.2.y1.49/32    192.168.y.17             0    100      0 65042 ?
*>i10.2.y2.49/32    192.168.y.33             0    100      0 65042 ?
*> 10.10.10.49/32   0.0.0.0                  0         32768 ?
P1>
```

Note	In these results, *x* is your SP number, and *y* is the number of the other SP. You need only check for the *x* routes.

Lab 7-1: Establishing Central Site Internet Connectivity with an MPLS VPN

Internet connectivity in MPLS VPN-based networks can be achieved through a dedicated Internet VPN. The dedicated Internet VPN approach provides better security because it completely isolates the service provider core (P routers) from the Internet. On the other hand, this approach is also less scalable; for example, you cannot export full Internet routing updates into an Internet VPN.

Complete this lab activity to practice what you learned in the related module.

Activity Objective

In this activity, you will remove the overlapping VPNS between customer A and customer B and then configure a separate VPN for Internet access for each customer. After completing this activity, you will be able to meet these objectives:

- Restore a simple customer VPN configuration
- Establish CE-PE connectivity for central site Internet access
- Establish central site CE-PE connectivity for Internet access across a separate MPLS VPN
- Establish remote site Internet connectivity through the central site router

Visual Objective

The figure illustrates what you will accomplish in this activity. For clarity, customer A connectivity is shown on the left side of the diagram, and customer B connectivity is shown on the right side of the diagram. Customer A is connected through POP *x*1, customer B is connected through POP *x*2.

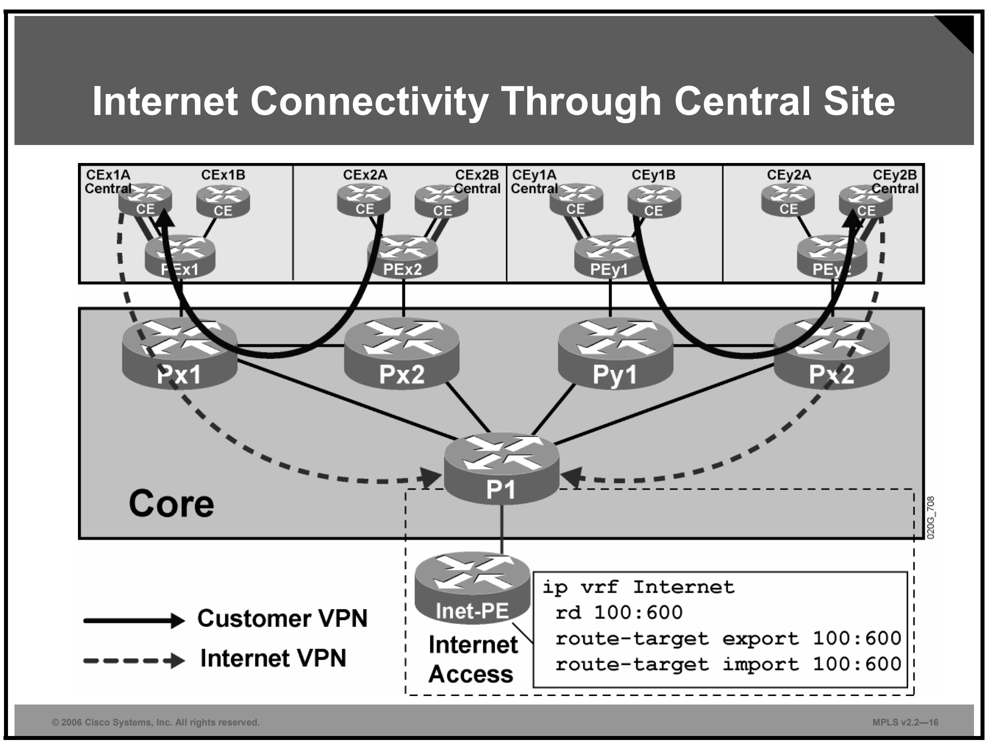

Note	The simulated Internet gateway router with an Internet VRF is preconfigured on P1.

You will first remove the overlapping VPNs between customer A and customer B, and reestablish simple VPN connectivity for each customer. You will then configure additional virtual links between the central site CE routers (CE*x*1A and CE*x*2B) and their PE routers. These separate circuits will connect from the central CE routers to their PE router to carry the Internet traffic.

The figure illustrates the new subinterfaces.

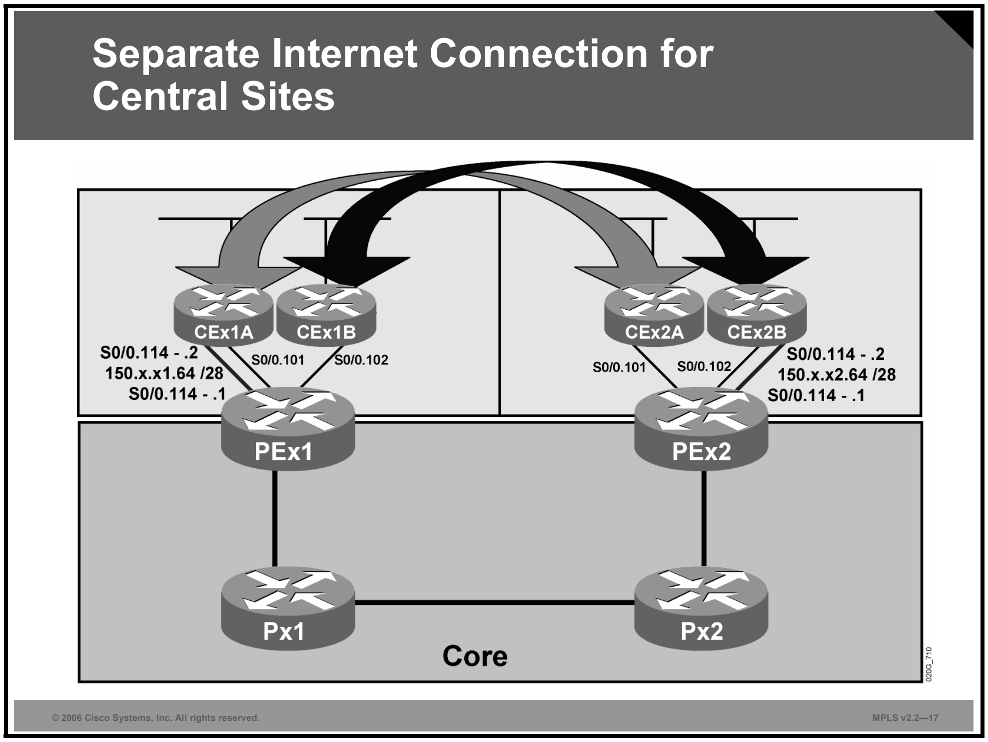

You will next create connectivity on the PE router between the Internet VPN and the customer central site for all Internet traffic in the SP. Each POP will be responsible for performing the configuration tasks on its PE router. The PE router will send only the default route from the Internet gateway to the central CE router.

Because the remote sites (CE*x*1B and CE*x*2A) will access the Internet using the MPLS VPN connection to their respective central sites, you will create default routes for each VPN pointing to the central CE router.

Note	Internet traffic within the SP will not go to the Internet gateway but will be appropriately routed by the PE routers.

Required Resources

This is the resource that is required to complete this activity:

- Cisco IOS documentation

Command List

The table describes some commands that are used in this activity. All other commands used in this lab have been described in previous labs.

Command	Description
ip prefix-list *name* **permit** *address mask*	Creates an IP prefix list that matches a specific address and mask
neighbor *ip-address* **prefix-list** *prefix-list-name* **out**	Filters BGP advertisements to a neighbor based on the prefix list
ping *host* **source** [*ip-address* | *interface-type interface-number*]	Pings a host using a specified source address

Task 1: Restore a Simple Customer VPN Configuration

In this task, you will remove the overlapping VRF CustAB. You will restore all customer A sites to VRF CustA and all customer B sites to VRF CustB with the following connectivity goals:

- CE*x*1A and CE*x*2A can communicate with each other, but not to CE*x**B devices.
- CE*x*1B and CE*x*2B can communicate with each other, but not to CE*x**A devices.

Activity Procedure

Complete these steps:

Step 1 Remove the address family BGP neighbor relationship between CE*x*1A and CE*x*2B on their respective PE router.

Step 2 Configure CE*x*1A with the CustA address family BGP neighbor relationship.

Step 3 Configure CE*x*2B with the CustB address family BGP neighbor relationship.

Activity Verification

You have completed this task when you attain these results:

- You have verified that the appropriate neighbor relationships are in place. Use the **show ip bgp** command to verify this.

```
PEx1#sh ip bgp vpn vrf CustA
BGP table version is 56, local router ID is 192.168.4.17
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: x:10 (default for vrf CustA)
*> 10.1.x1.16/28    150.x.x1.17              0             0 65041 ?
*> 10.1.x1.49/32    150.x.x1.17              0             0 65041 ?
*>i10.1.x2.16/28    192.168.x.33             0    100      0 65041 ?
*                   150.x.x1.49            200             0 65041 ?
*>i10.1.x2.49/32    192.168.x.33             0    100      0 65041 ?
*                   150.x.x1.49            200             0 65041 ?
*>i10.10.10.49/32   192.168.100.129          0    100      0 ?
r> 150.x.x1.16/28   150.x.x1.17              0             0 65041 ?
r>i150.x.x1.48/28   192.168.x.33             0    100      0 65041 ?
r                   150.x.x1.49            200             0 65041 ?
*>i150.x.x2.16/28   192.168.x.33             0    100      0 65041 ?
*                   150.x.x1.49            200             0 65041 ?
PEx1#

PEx2#sh ip bgp vpn vrf CustB
BGP table version is 56, local router ID is 192.168.x.33
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: x:20 (default for vrf CustB)
*>i10.2.x1.16/28    192.168.x.17             0    100      0 65042 ?
*                   150.x.x2.49            200             0 65042 ?
*>i10.2.x1.49/32    192.168.x.17             0    100      0 65042 ?
*                   150.x.x2.49            200             0 65042 ?
*> 10.2.x2.16/28    150.x.x2.33              0             0 65042 ?
*> 10.2.x2.49/32    150.x.x2.33              0             0 65042 ?
*>i10.10.10.49/32   192.168.100.129          0    100      0 ?
*>i150.x.x1.32/28   192.168.x.17             0    100      0 65042 ?
*                   150.x.x2.49            200             0 65042 ?
r> 150.x.x2.32/28   150.x.x2.33              0             0 65042 ?
```

```
r>i150.x.x2.48/28   192.168.x.17             0    100      0 65042 ?
r                   150.x.x2.49            200             0 65042 ?
PEx2#
```

- From your customer router, you have verified that you can ping the loopback interface of the remote customer router.

```
CEx1A#ping 10.1.x2.49

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.x2.49, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 148/148/149 ms


CEx2B#ping 10.2.x1.49

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.2.x1.49, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 148/148/149 ms
```

Note	You may need to issue a **clear ip bgp *** command on the CE router to force the propagation of the new routes.

Task 2: Establish CE-PE Connectivity for Central Site Internet Access

In this task, you will add a new subinterface to support the Internet VPN on the central site router.

Activity Procedure

Complete these steps:

Step 1 Create a separate subinterface (S0/0.114) on the central router of the customer using the address information in the table.

Router ID	IP Address	DLCI
CE*x*1A	150.*x*.*x*1.66/28	114
CE*x*2B	150.*x*.*x*2.66/28	114

Step 2 Activate the new interface in the IGP routing process and make the interface passive.

Step 3 Create a separate subinterface (S0/0.114) on the PE routers using the address information in this table.

Router ID	IP Address	DLCI
PE*x*1	150.*x*.*x*1.65/28	114
PE*x*2	150.*x*.*x*2.65/28	114

Step 4 Activate the new interface in the IGP routing process and make the interface passive.

Note Global routing between your PE router and P1 was established in Lab 6-2, "Merging Service Providers."

Activity Verification

You have completed this task when you attain these results:

- You have used the **show ip interface** command to verify the status of the new interfaces.

```
CEx1A#show ip interface s0/0.114
Serial0/0.114 is up, line protocol is up
  Internet address is 150.x.x1.66/28
  Broadcast address is 255.255.255.255
  Address determined by setup command
  MTU is 1500 bytes
  ************** output omitted ************************

PEx1#show ip interface s0/0.114
Serial0/0.114 is up, line protocol is up
```



```
Internet address is 150.x.x1.65/28
Broadcast address is 255.255.255.255
Address determined by setup command
MTU is 1500 bytes
************** output omitted ************************
```

Task 3: Establish Central Site Connectivity for Internet Access

Another group in the service provider organization has already created a VPN in the provider core to carry the Internet traffic. You will connect the central customer site CE to the Internet VPN to support all Internet traffic from the customer. Each POP will be responsible for performing the configuration tasks on its PE router.

To control Internet routes forwarded to the customer, the PE router will send only the default route from the Internet gateway to the central CE router.

Activity Procedure

Complete these steps:

Step 1 Create /24 summary routes for both customer sites with default routes to the Null0 interface on the central CE router.

Step 2 Add the appropriate PE router neighbor statement to the BGP process on the CE router.

Step 3 Add the summary networks to the BGP routing process on the central CE router.

Step 4 Create a VRF named "Internet" on the PE router. Use RD 100:600, and import RT 100:600.

Step 5 Create a route-map to export only the default route and the /24 summary routes from the customer.

Step 6 Place the new interface (Se0/0.114) that will support the central site CE router (CE*x*1A or CE*x*2B) into the Internet VRF on the PE router.

Step 7 Create a prefix list to permit only the default route.

Step 8 Add the central site router neighbor statements to the IPv4 VRF address family for the Internet VRF on the PE router.

Activity Verification

You have completed this task when you attain these results:

- You have verified that the central site CE router is receiving only the default route across the Internet subinterface from its PE neighbor and that /24 summary routes for both customer sites are in the routing table.

```
CEx1A#sh ip route
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, 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 150.x.x1.65 to network 0.0.0.0

     10.0.0.0/8 is variably subnetted, 7 subnets, 3 masks
B       10.1.x2.49/32 [20/0] via 150.x.x1.18, 01:27:13
C       10.1.x1.49/32 is directly connected, Loopback0
S       10.1.x2.0/24 is directly connected, Null0
S       10.1.x1.0/24 is directly connected, Null0
B       10.1.x2.16/28 [20/0] via 150.x.x1.18, 01:27:13
C       10.1.x1.16/28 is directly connected, Ethernet0/0
B       10.10.10.49/32 [20/0] via 150.x.x1.18, 00:35:11
     150.x.0.0/28 is subnetted, 4 subnets
B       150.x.x1.48 [20/0] via 150.x.x1.18, 01:27:14
B       150.x.x2.16 [20/0] via 150.x.x1.18, 01:27:14
C       150.x.x1.16 is directly connected, Serial0/0.101
C       150.x.x1.64 is directly connected, Serial0/0.114
B*   0.0.0.0/0 [20/0] via 150.x.x1.65, 00:34:53
CEx1A#

CEx2B#show ip route
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, 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 150.x.x2.65 to network 0.0.0.0

     10.0.0.0/8 is variably subnetted, 7 subnets, 3 masks
B       10.2.x1.49/32 [20/0] via 150.x.x2.34, 00:35:24
C       10.2.x2.49/32 is directly connected, Loopback0
S       10.2.x1.0/24 is directly connected, Null0
S       10.2.x2.0/24 is directly connected, Null0
B       10.2.x1.16/28 [20/0] via 150.x.x2.34, 00:35:24
C       10.2.x2.16/28 is directly connected, Ethernet0/0
B       10.10.10.49/32 [20/0] via 150.x.x2.34, 00:35:24
     150.x.0.0/28 is subnetted, 4 subnets
B       150.x.x2.48 [20/0] via 150.x.x2.34, 00:35:25
C       150.x.x2.32 is directly connected, Serial0/0.102
B       150.x.x1.32 [20/0] via 150.x.x2.34, 00:35:25
```

```
C       150.x.x2.64 is directly connected, Serial0/0.114
B*   0.0.0.0/0 [20/0] via 150.x.x2.65, 00:15:17
CEx2B#
```

- You have verified that the PE router is receiving Internet routes from the central CE router as well as the Internet gateway. The /24 summary routes for both customer sites are in the Internet routing table.

```
PEx1#show ip route vrf Internet

Routing Table: Internet
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 192.168.100.129 to network 0.0.0.0

     10.0.0.0/8 is variably subnetted, 10 subnets, 3 masks
B       10.1.x1.49/32 [20/0] via 150.x.x1.66, 00:30:46
B       10.1.y1.0/24 [200/0] via 192.168.y.17, 00:24:37
B       10.2.y1.0/24 [200/0] via 192.168.y.33, 00:24:37
B       10.2.x1.0/24 [200/0] via 192.168.x.33, 00:16:51
B       10.1.x2.0/24 [20/0] via 150.x.x1.66, 00:20:58
B       10.2.x2.0/24 [200/0] via 192.168.x.33, 00:16:51
B       10.1.x1.0/24 [20/0] via 150.x.x1.66, 00:21:27
B       10.2.y2.0/24 [200/0] via 192.168.y.33, 00:24:38
B       10.1.y2.0/24 [200/0] via 192.168.y.17, 00:24:38
B       10.1.x1.16/28 [20/0] via 150.x.x1.66, 00:30:47
     150.x.0.0/28 is subnetted, 2 subnets
B       150.x.x1.16 [20/0] via 150.4.41.66, 00:3:48
C       150.x.x1.64 is directly connected, Serial0/0.114
B*   0.0.0.0/0 [200/0] via 192.168.100.129, 00:16:52
PEx1#
```

Note	In these results, *x* is your SP number, and *y* is the number of the other SP in the merged SP network. Depending on the status of the other groups, you could see the /24 subnets for CE*x*1A, CE*x*2A, CE*x*1B, CE*x*2B, CE*y*1A, CE*y*2A, CE*y*1B, and CE*y*2B.

```
PEx2#show ip route vrf Internet

Routing Table: Internet
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 192.168.100.129 to network 0.0.0.0

     10.0.0.0/8 is variably subnetted, 10 subnets, 3 masks
B       10.2.x2.49/32 [20/0] via 150.x.x2.66, 00:51:34
B       10.1.31.0/24 [200/0] via 192.168.x.17, 00:51:27
B       10.2.31.0/24 [200/0] via 192.168.x.33, 00:51:27
B       10.2.x1.0/24 [20/0] via 150.x.x2.66, 00:51:34
B       10.1.x2.0/24 [200/0] via 192.168.x.17, 00:51:27
B       10.2.x2.0/24 [20/0] via 150.x.x2.66, 00:51:34
B       10.1.x1.0/24 [200/0] via 192.168.x.17, 00:51:27
B       10.2.32.0/24 [200/0] via 192.168.x.33, 00:51:28
B       10.1.32.0/24 [200/0] via 192.168.x.17, 00:51:28
B       10.2.x2.16/28 [20/0] via 150.x.x2.66, 00:51:35
     150.x.0.0/28 is subnetted, 2 subnets
B       150.x.x2.32 [20/0] via 150.x.x2.66, 00:51:37
C       150.x.x2.64 is directly connected, Serial0/0.114
B*   0.0.0.0/0 [200/0] via 192.168.100.129, 00:31:14
PEx2#
```

Note In these results, *x* is your SP number, and *y* is the number of the other SP in the merged SP network. Depending on the status of the other groups, you could see the /24 subnets for CE*x*1A, CE*x*2A, CE*x*1B, CE*x*2B, CE*y*1A, CE*y*2A, CE*y*1B, and CE*y*2B.

- You have verified that the /24 summary routes for both customer sites are in the Internet routing table on the P1 router.

```
P1>show ip route vrf Internet

Routing Table: Internet
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 10.0.0.2 to network 0.0.0.0

     172.16.0.0/24 is subnetted, 1 subnets
C       172.16.0.0 is directly connected, Loopback172
     10.0.0.0/8 is variably subnetted, 16 subnets, 2 masks
C       10.0.0.0/24 is directly connected, Loopback200
B       10.1.x1.0/24 [200/0] via 192.168.x.17, 00:28:34
B       10.2.x1.0/24 [200/0] via 192.168.x.33, 00:28:34
B       10.2.y1.0/24 [200/0] via 192.168.y.33, 01:00:49
B       10.1.y2.0/24 [200/0] via 192.168.y.17, 00:24:34
B       10.2.y2.0/24 [200/0] via 192.168.y.33, 00:00:50
B       10.1.y1.0/24 [200/0] via 192.168.y.17, 00:25:20
B       10.2.x2.0/24 [200/0] via 192.168.x.33, 00:28:35
B       10.1.x2.0/24 [200/0] via 192.168.x.17, 00:28:37
C       10.2.72.0/24 is directly connected, Loopback272
C       10.1.71.0/24 is directly connected, Loopback271
C       10.1.91.0/24 is directly connected, Loopback291
C       10.2.92.0/24 is directly connected, Loopback292
C       10.2.82.0/24 is directly connected, Loopback282
C       10.1.81.0/24 is directly connected, Loopback281
C    192.168.0.0/24 is directly connected, Loopback192
S*   0.0.0.0/0 [1/0] via 10.0.0.2
P1>
```

Note In these results, *x* is your SP number, and *y* is the number of the other SP in the merged SP network. Depending on the status of the other groups, you could see the /24 subnets for CE*x*1A, CE*x*2A, CE*x*1B, CE*x*2B, CE*y*1A, CE*y*2A, CE*y*1B, and CE*y*2B.

P1 also has loopbacks simulating Internet devices. You should be able to ping 10.0.0.1, 172.176.0.1, and 192.168.0.1 from any prefix in this routing table.

- You have verified that the central customer site devices can reach the simulated networks the P1 Internet VRF routing table.

```
CEx1A#ping 172.16.0.1 source loopback0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.0.1, timeout is 2 seconds:
Packet sent with a source address of 10.1.x1.49
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 88/93/112 ms
CEx1A#



CEx2B#ping 10.0.0.1 source loopback0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.0.0.1, timeout is 2 seconds:
Packet sent with a source address of 10.2.x2.49
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 88/117/237 ms
CEx2B#trace
Protocol [ip]:
Target IP address: 10.0.0.1
Source address: 10.2.x2.17
Numeric display [n]:
Timeout in seconds [3]:
Probe count [3]:
Minimum Time to Live [1]:
Maximum Time to Live [30]:
Port Number [33434]:
Loose, Strict, Record, Timestamp, Verbose[none]:
Type escape sequence to abort.
Tracing the route to 10.0.0.1

  1 150.x.x2.65 17 msec 16 msec 16 msec
  2 10.0.0.1 [AS 65001] 44 msec *  44 msec
CEx2B#
```

Note	The packet from the central CE router should go to its PE router (150.*x.x*1.65) through the Internet interface to reach the Internet address.

Task 4: Establish Remote Site Internet Connectivity Through the Central Site Router

Because the remote sites (CE*x*1B and CE*x*2A) will access the Internet using the MPLS VPN connection to their respective central sites, you will create default routes for each VPN pointing to the central CE router.

Activity Procedure

Complete these steps:

Step 1 On your PE router, add a default VRF route pointing to Loopback0 on the central CE router.

Step 2 Add the **redistribute static** command and the **default-information originate** command to your address family to propagate this route to the CEs.

Activity Verification

You have completed this task when you attain these results:

- You have verified that the remote CE router has a default route across the MPLS VPN to its CE neighbor.

```
CEx1B#show ip route
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, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route

Gateway of last resort is 10.2.x2.49 to network 0.0.0.0

     10.0.0.0/8 is variably subnetted, 7 subnets, 3 masks
C       10.2.x1.49/32 is directly connected, Loopback0
B       10.2.x2.49/32 [20/0] via 150.x.x1.34, 00:02:21
B       10.2.x1.0/24 [20/0] via 150.x.x1.34, 00:02:21
B       10.2.x2.0/24 [20/0] via 150.x.x1.34, 00:02:21
C       10.2.x1.16/28 is directly connected, Ethernet0/0
B       10.2.x2.16/28 [20/0] via 150.x.x1.34, 00:02:21
B       10.10.10.49/32 [20/0] via 150.x.x1.34, 00:22:09
     150.x.0.0/28 is subnetted, 4 subnets
C       150.x.x2.48 is directly connected, Serial0/0.113
B       150.x.x2.32 [20/0] via 150.x.x1.34, 00:02:22
C       150.x.x1.32 is directly connected, Serial0/0.102
B       150.x.x2.64 [20/0] via 150.x.x1.34, 00:02:22
B*   0.0.0.0/0 [1/0] via 10.2.x2.49
CEx1B#
```

```
CEx2A#show ip route
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, 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 10.1.x1.49 to network 0.0.0.0

     10.0.0.0/8 is variably subnetted, 7 subnets, 3 masks
C       10.1.x2.49/32 is directly connected, Loopback0
B       10.1.x1.49/32 [20/0] via 150.x.x2.18, 00:08:00
B       10.1.x2.0/24 [20/0] via 150.x.x2.18, 00:08:00
B       10.1.x1.0/24 [20/0] via 150.x.x2.18, 00:08:00
C       10.1.x2.16/28 is directly connected, Ethernet0/0
B       10.1.x1.16/28 [20/0] via 150.x.x2.18, 00:08:00
B       10.10.10.49/32 [20/0] via 150.x.x2.18, 00:08:00
     150.x.0.0/28 is subnetted, 4 subnets
C       150.x.x1.48 is directly connected, Serial0/0.113
C       150.x.x2.16 is directly connected, Serial0/0.101
B       150.x.x1.16 [20/0] via 150.x.x2.18, 00:08:02
B       150.x.x1.64 [20/0] via 150.x.x2.18, 00:08:02
S*   0.0.0.0/0 [1/0] via 10.1.x1.49
CEx2A#
```

- You have verified that the PE routers have the appropriate default routes across the MPLS VPN to the CE neighbor.

```
PEx1#sh ip ro vrf CustB

Routing Table: CustB
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 192.168.x.33 to network 0.0.0.0

     10.0.0.0/8 is variably subnetted, 7 subnets, 3 masks
B       10.2.x1.49/32 [20/0] via 150.x.x1.33, 00:14:37
B       10.2.x2.49/32 [200/0] via 192.168.x.33, 00:09:35
B       10.2.x1.16/28 [20/0] via 150.x.x1.33, 00:14:37
B       10.2.x2.16/28 [200/0] via 192.168.x.33, 00:09:35
```



```
B       10.10.10.49/32 [200/0] via 192.168.100.129, 00:14:22
B       10.2.x1.0/24 [200/0] via 192.168.x.33, 00:11:05
B       10.2.x2.0/24 [200/0] via 192.168.x.33, 00:11::05
     150.x.0.0/28 is subnetted, 4 subnets
C       150.x.x1.32 is directly connected, Serial0/0.102
B       150.x.x2.32 [200/0] via 192.168.x.33, 00:09:36
B       150.x.x2.48 [20/0] via 150.x.x1.33, 00:16:13
B       150.x.x2.64 [200/0] via 192.168.x.33, 00:09:40
B*   0.0.0.0/0 [200/0] via 192.168.x.33, 00:09:40
PEx1#

PEx2#sh ip ro vrf CustA

Routing Table: CustA
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 192.168.x.17 to network 0.0.0.0

     10.0.0.0/8 is variably subnetted, 7 subnets, 3 masks
B       10.1.x1.49/32 [200/0] via 192.168.x.17, 00:18:15
B       10.1.x2.49/32 [20/0] via 150.x.x2.17, 00:15:46
B       10.1.x1.16/28 [200/0] via 192.168.x.17, 00:18:15
B       10.1.x2.16/28 [20/0] via 150.x.x2.17, 00:15:46
B       10.1.x1.0/24 [200/0] via 192.168.x.17, 00:18:15
B       10.10.10.49/32 [200/0] via 192.168.100.129, 00:18:15
B       10.1.x2.0/24 [200/0] via 192.168.x.17, 00:18:15
     150.x.0.0/28 is subnetted, 4 subnets
B       150.x.x1.48 [20/0] via 150.x.x2.17, 00:15:47
B       150.x.x1.16 [200/0] via 192.168.x.17, 00:18:16
C       150.x.x2.16 is directly connected, Serial0/0.101
B       150.x.x1.64 [200/0] via 192.168.x.17, 00:18:17
B*   0.0.0.0/0 [1/0] via 192.168.x.17, 00:18:17
PEx2#
```

- You have used the **ping** and **trace** commands to verify that you can reach remote devices.

```
CEx2A#ping 10.0.0.1 source loopback0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.0.0.1, timeout is 2 seconds:
```

```
Packet sent with a source address of 10.1.x2.49
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 232/249/317 ms
      CEx1B#ping 172.16.0.1 source loopback0

      Type escape sequence to abort.
      Sending 5, 100-byte ICMP Echos to 10.1.71.149, timeout is 2 seconds:
      Packet sent with a source address of 10.1.x2.49
      !!!!!
      Success rate is 100 percent (5/5), round-trip min/avg/max = 232/249/317 ms
      CEx1B#

      CEx2A#trace
      Protocol [ip]:
      Target IP address: 10.0.0.1
      Source address: 10.1.x2.17
      Numeric display [n]:
      Timeout in seconds [3]:
      Probe count [3]:
      Minimum Time to Live [1]:
      Maximum Time to Live [30]:
      Port Number [33434]:
      Loose, Strict, Record, Timestamp, Verbose[none]:
      Type escape sequence to abort.
      Tracing the route to 10.0.0.1

        1 150.x.x2.18 12 msec 12 msec 17 msec
        2 150.x.x1.18 [AS 65001] [MPLS: Label 39 Exp 0] 120 msec 116 msec 116 msec
        3 150.x.x1.17 [AS 65001] 76 msec 76 msec 76 msec
        4 150.x.x1.65 [AS 65001] 89 msec 176 msec 260 msec
        5 10.0.0.1 121 msec *  117 msec
      CEx2A#
```

Note If you trace the path, the packet from the remote CE router should go to its PE (150.*x*.*x*2.18), then to the central site PE (150.*x*.*x*1.18 outbound), then to the central site CE (150.*x*.*x*1.17) , and then out the central site CE Internet interface to the central site PE (150.*x*.*x*1.65) to reach the Internet address.

Lab 8-1: Implementing Basic MPLS TE

In this exercise, you will establish MPLS traffic tunnels to support a new requirement of the SP management.

Complete this lab activity to practice what you learned in the related module.

Activity Objective

In this activity, you will configure MPLS TE to forward all traffic through the P1 router. The links between P*x*1 and P*x*2 will be retained for backup purposes. You will use MPLS TE to direct customer traffic across the P1 router. After completing this activity, you will be able to meet these objectives:

- Follow traffic flows in a MPLS backbone
- Configure basic MPLS TE support on your routers
- Configure MPLS TE tunnels

Visual Objective

The figure illustrates what you will accomplish in this activity.

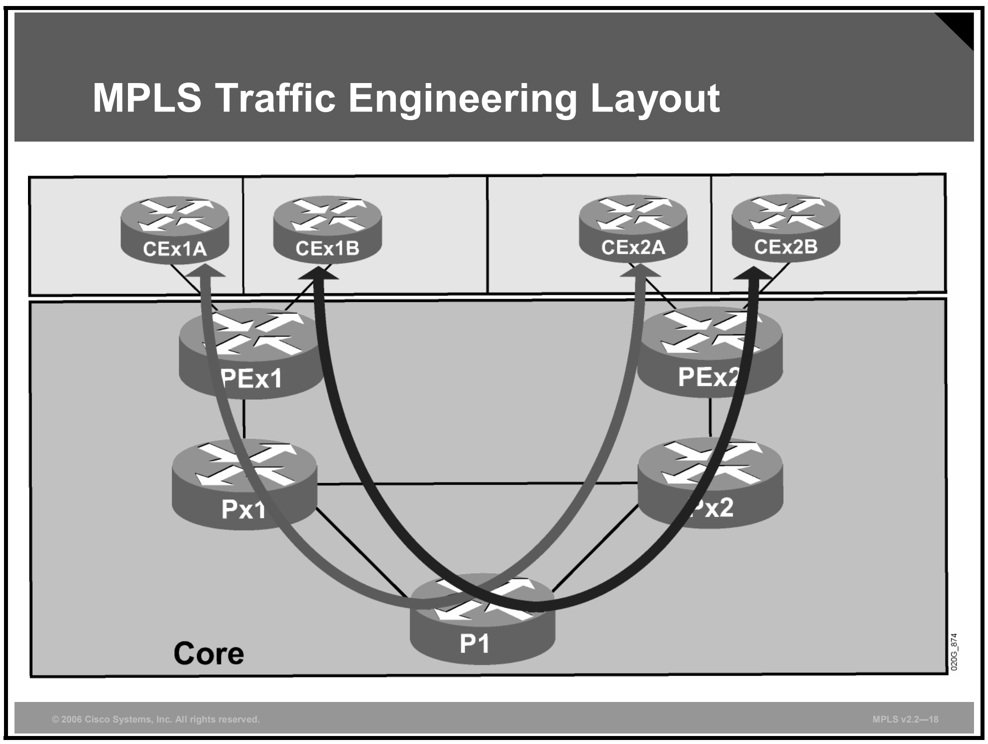

Required Resources

This is the resource that is required to complete this activity:

- Cisco IOS documentation

Command List

The table describes the commands that are used in this activity.

Command	Description
interface tunnel *number*	Creates a tunnel interface
ip explicict-path name *name*	Enters the subcommand mode for IP explicit paths and creates or modifies the specified path
ip rsvp bandwidth [*max-reservable* [*single-request*]]	Enables RSVP on an interface
metric-style wide	Specifies the IS-IS metric type
mpls ip	Starts MPLS support and LDP on an interface
mpls ip propagate-ttl	Enables MPLS TTL propagation support
mpls traffic-eng *level*	Configures IS-IS TE support
mpls traffic-eng logging	Configures MPLS TE logging
mpls traffic-eng reoptimize timers frequency 30	Controls the frequency with which tunnels with established LSPs are checked for better LSPs
mpls traffic-eng router-id *interface*	Defines the IP address that is used as the tunnel destination for this router
mpls traffic-eng tunnels	Enables MPLS TE globally or on an interface
next-address *ip-address*	Specifies the next IP address in the explicit path
router isis	Starts IS-IS configuration
show mpls traffic-eng autoroute	Displays autoroute information
show mpls traffic-eng tunnels	Displays detailed tunnel information
show mpls traffic-eng tunnels brief	Displays the overall status of MPLS TE tunnels
tunnel destination *ip-address*	Specifies the tunnel tailend router
tunnel mode mpls traffic-eng	Specifies tunnel encapsulation mode
tunnel mpls traffic-eng autoroute announce	Enables autoroute on an MPLS TE tunnel
tunnel mpls traffic-eng bandwidth *bandwidth*	Specifies the bandwidth that is reserved by the MPLS TE tunnel
tunnel mpls traffic-eng path-option *number* **explicit name** *path-name*	Specifies that the MPLS TE path option will use an explicit path
tunnel mpls traffic-eng priority *setup hold*	Specifies MPLS TE tunnel setup and hold priority

Task 1: Log the Existing Traffic Flow

In this task, you will log the existing traffic flow pattern to establish the current path through the network.

Exercise Procedure

Complete these steps:

Step 1 Enable MPLS TTL propagation on all PE and P routers.

Step 2 Perform a traceroute between your two assigned CE routers.

Exercise Verification

You have completed this task when you attain these results:

- You have used the **trace** command to verify the route between the devices.

```
CEx1B#trace 10.2.x2.49

Type escape sequence to abort.
Tracing the route to 10.2.x2.49

  1 150.x.x1.34 12 msec 12 msec 12 msec
  2 192.168.x.50 [MPLS: Labels 21/49 Exp 0] 397 msec 256 msec 301 msec
  3 192.168.x.114 [MPLS: Labels 17/49 Exp 0] 204 msec 196 msec 197 msec
  4 150.x.x2.34 [AS 65001] [MPLS: Label 49 Exp 0] 116 msec 116 msec 116 msec
  5 150.x.x2.33 [AS 65001] 76 msec *  72 msec
CEx1B#



CEx2B#trace 10.2.x1.49

Type escape sequence to abort.
Tracing the route to 10.2.x1.49

  1 150.x.x2.34 16 msec 16 msec 12 msec
  2 192.168.x.66 [AS 65001] [MPLS: Labels 23/42 Exp 0] 297 msec 300 msec 297 msec
  3 192.168.x.113 [AS 65001] [MPLS: Labels 18/42 Exp 0] 244 msec 193 msec 196 msec
  4 150.x.x1.34 [AS 65001] [MPLS: Label 42 Exp 0] 116 msec 116 msec 117 msec
  5 150.x.x1.33 [AS 65001] 76 msec *  72 msec
CEx2B#
```

Task 2: Configure MPLS TE Support on the PE and P Routers

In this task, you will enable global features that will allow TE tunnels to be built.

Exercise Procedure

Complete these steps:

Step 1 Configure global MPLS TE support.

Step 2 Configure IS-IS support to generate and accept only new-style TLV objects. Configure IS-IS to support level 2 TE.

Step 3 Use "loopback0" as the router ID on the P*x*1, P*x*2, PE*x*1, and PE*x*2 routers.

Step 4 Configure MPLS TE support and reservable bandwidth on all links between the backbone routers. Allow up to 128 kbps of reservable bandwidth on any single subinterface and up to 128 kbps of reservable bandwidth on physical interfaces.

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

Exercise Verification

You have completed this task when you attain these results:

- You have used the **show ip rsvp interface** command to confirm the RVSP information that has been configured. Verify that the proper interface and bandwidths have been enabled.

```
PEx1#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

PEx2#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

Px1#show ip rsvp interface
interface    allocated  i/f max  flow max sub max
Se0/0        0          128K     128K     0
Se0/0.112    0          128K     128K     0
Se0/0.2x1    0          128K     128K     0
Se0/0.111    0          128K     128K     0
```

```
Px2#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
Se0/0.2x2     0           128K      128K     0
```

Task 3: Configure MPLS TE Tunnels

In this task, you will create the traffic tunnels between PE routers.

Exercise Procedure

Complete these steps:

Step 1 On both of your PE routers, configure an MPLS TE tunnel toward the other PE router in your SP passing through the P1 router (for example, PE*x*1 → P*x*1 → P1 → P*x*2 → PE*x*2). Use the tunnel parameters shown in the table.

Parameter	Value
IP address	Loopback 0
Tunnel bandwidth	100 kbps
Path options	Explicit path selection
Tunnel priority	Setup = 0, hold = 0
Autoroute	Enabled

Step 2 Enable MPLS forwarding on the MPLS TE tunnels to establish end-to-end LSPs to be used for MPLS VPN traffic.

Step 3 Force faster tunnel optimization with the global command **mpls traffic-eng reoptimize timers frequency 30**.

Exercise Verification

You have completed this task when you attain these results:

- You have used the **show mpls traffic-engineering tunnels brief** command to display tunnels that are going through a particular router and verified the operational state of the tunnels that are originating in the router. You should get a printout similar to this example.

```
PEx1#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
PEx1_t0                          192.168.x.33     -         Se0/0.11  up/up
PEx2_t0                          192.168.x.17     Se0/0.11  -         up/up
Displayed 1 (of 1) heads, 0 (of 0) midpoints, 1 (of 1) tailS
```

```
PEx2#show mpls traffic-eng tunnels brief
Signaling Summary:
    LSP Tunnels Process:            running
    RSVP Process:                   running
    Forwarding:                     enabled
    Periodic reoptimization:        every 30 seconds, next in 2 seconds
    Periodic auto-bw collection:    disabled
TUNNEL NAME                      DESTINATION      UP IF     DOWN IF
STATE/PROT
PEx2_t0                          192.168.x.17     -         Se0/0.11  up/up
PEx1_t0                          192.168.x.33     Se0/0.11  -         up/up
Displayed 1 (of 1) heads, 0 (of 0) midpoints, 1 (of 1) tails
```

Note If the tunnel does not come up, check these conditions:

a) The required global commands are implemented on the P and PE routers.
b) The required ISIS commands are implemented on the P and PE routers.
c) The required interface commands are implemented on the P and PE routers.
d) The explicit tunnel path is correctly configured.
e) The ingress tunnel interface commands are correctly configured.

- You have used **show ip route** to verify that the tunnel is in the global routing table.

```
PEx1#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, 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

     192.168.x.0/24 is variably subnetted, 7 subnets, 2 masks
i L2    192.168.x.97/32 [115/30] via 192.168.x.50, Serial0/0.111
i L2    192.168.x.112/28 [115/20] via 192.168.x.50, Serial0/0.111
i L2    192.168.x.64/28 [115/30] via 192.168.x.50, Serial0/0.111
i L2    192.168.x.81/32 [115/20] via 192.168.x.50, Serial0/0.111
i L2    192.168.x.33/32 [115/40] via 192.168.x.33, Tunnel0
C       192.168.x.48/28 is directly connected, Serial0/0.111
C       192.168.x.17/32 is directly connected, Loopback0
     192.168.100.0/24 is variably subnetted, 5 subnets, 2 masks
i L2    192.168.100.8/29 [115/30] via 192.168.x.50, Serial0/0.111
i L2    192.168.100.24/29 [115/20] via 192.168.x.50, Serial0/0.111
i L2    192.168.100.16/29 [115/30] via 192.168.x.50, Serial0/0.111
i L2    192.168.100.32/29 [115/30] via 192.168.x.50, Serial0/0.111
i L2    192.168.100.129/32 [115/30] via 192.168.x.50, Serial0/0.111
     192.168.y.0/24 is variably subnetted, 7 subnets, 2 masks
i L2    192.168.y.97/32 [115/40] via 192.168.x.50, Serial0/0.111
i L2    192.168.y.112/28 [115/40] via 192.168.x.50, Serial0/0.111
i L2    192.168.y.64/28 [115/40] via 192.168.x.50, Serial0/0.111
i L2    192.168.y.81/32 [115/40] via 192.168.x.50, Serial0/0.111
i L2    192.168.y.33/32 [115/50] via 192.168.x.50, Serial0/0.111
i L2    192.168.y.48/28 [115/40] via 192.168.x.50, Serial0/0.111
i L2    192.168.y.17/32 [115/50] via 192.168.x.50, Serial0/0.111
PEx1#
```

Note In these results, *x* is your SP, and *y* is the other SP as viewed from PE*x*1. Depending on how far along the other SP is in configuring its devices, you may not see the *y* routes. You need be concerned only about the *x* routes.

```
PEx2#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, 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

     192.168.x.0/24 is variably subnetted, 7 subnets, 2 masks
i L2    192.168.x.97/32 [115/20] via 192.168.x.66, Serial0/0.111
i L2    192.168.x.112/28 [115/20] via 192.168.x.66, Serial0/0.111
C       192.168.x.64/28 is directly connected, Serial0/0.111
i L2    192.168.x.81/32 [115/30] via 192.168.x.66, Serial0/0.111
C       192.168.x.33/32 is directly connected, Loopback0
i L2    192.168.x.48/28 [115/30] via 192.168.x.66, Serial0/0.111
i L2    192.168.x.17/32 [115/40] via 192.168.x.17, Tunnel0
     192.168.100.0/24 is variably subnetted, 5 subnets, 2 masks
i L2    192.168.100.8/29 [115/30] via 192.168.x.66, Serial0/0.111
i L2    192.168.100.24/29 [115/30] via 192.168.x.66, Serial0/0.111
i L2    192.168.100.16/29 [115/30] via 192.168.x.66, Serial0/0.111
i L2    192.168.100.32/29 [115/20] via 192.168.x.66, Serial0/0.111
i L2    192.168.100.129/32 [115/30] via 192.168.x.66, Serial0/0.111
     192.168.y.0/24 is variably subnetted, 7 subnets, 2 masks
i L2    192.168.y.97/32 [115/40] via 192.168.x.66, Serial0/0.111
i L2    192.168.y.112/28 [115/40] via 192.168.x.66, Serial0/0.111
i L2    192.168.y.64/28 [115/40] via 192.168.x.66, Serial0/0.111
i L2    192.168.y.81/32 [115/40] via 192.168.x.66, Serial0/0.111
i L2    192.168.y.33/32 [115/50] via 192.168.x.66, Serial0/0.111
i L2    192.168.y.48/28 [115/40] via 192.168.x.66, Serial0/0.111
i L2    192.168.y.17/32 [115/50] via 192.168.x.66, Serial0/0.111
PEx2#
```

Note In these results, *x* is your SP, and *y* is the other SP as viewed from PE*x*2. Depending on how far along the other SP is in configuring its devices, you may not see the *y* routes. You need be concerned only about the "*x*" routes.

- You have used the **show ip cef** *network* command to verify that the IP traffic toward the tunnel destination gets forwarded through the tunnel interface.

```
PEx1#show ip cef 192.168.x.33
192.168.x.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.x.33, Tunnel0, 11 dependencies
    next hop 192.168.x.33, Tunnel0
    valid adjacency
    tag rewrite with Tu0, point2point, tags imposed: {33}

PEx2#show ip cef 192.168.x.17
192.168.x.17/32, version 27, epoch 0
0 packets, 0 bytes
  tag information set
    local tag: 32
    fast tag rewrite with Tu0, point2point, tags imposed: {31}
  via 192.168.x.17, Tunnel0, 11 dependencies
    next hop 192.168.x.17, Tunnel0
    valid adjacency
    tag rewrite with Tu0, point2point, tags imposed: {31}
```

- You have used the **show ip cef vrf** *name network* command to verify that the MPLS VPN traffic gets forwarded through the tunnel interface.

```
PEx1#show ip cef vrf CustA 10.1.x2.49
10.1.x2.49/32, 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.x.33, 0 dependencies, recursive
    next hop 192.168.x.33, Tunnel0 via 192.168.x.33/32
    valid adjacency
    tag rewrite with Tu0, point2point, tags imposed: {33 34}
```

```
PEx2#sh ip cef vrf CustA 10.1.x1.49
10.1.4x.49/32, version 12, epoch 0
0 packets, 0 bytes
  tag information set, shared
    local tag: VPN-route-head
    fast tag rewrite with Tu0, point2point, tags imposed: {31 35}
  via 192.168.x.17, 1 dependency, recursive
    next hop 192.168.x.17, Tunnel0 via 192.168.x.17/32
    valid adjacency
    tag rewrite with Tu0, point2point, tags imposed: {31 35}
```

- You have repeated the traces that you did in Task 1 to verify that the trace now bypasses the direct P*x*1-to-P*x*2 link and passes through the P router links.

```
CEx1A>trace 10.1.x2.49

Type escape sequence to abort.
Tracing the route to 10.1.x2.49

  1 150.x.x1.18 12 msec 12 msec 12 msec
  2 192.168.x.50 [AS 65001] [MPLS: Labels 33/34 Exp 0] 277 msec 300 msec 301 msec
  3 192.168.100.** [AS 65001] [MPLS: Labels 32/34 Exp 0] 368 msec 385 msec 292 msec
  4 192.168.100.** [AS 65001] [MPLS: Labels 33/34 Exp 0] 217 msec 324 msec 273 msec
  5 150.x.x2.18 [AS 65001] [MPLS: Label 34 Exp 0] 144 msec 148 msec 144 msec
  6 150.x.x2.17 [AS 65001] 92 msec *  177 msec
```

Note	The third and fourth router addresses will vary, based on the SP.

```
CEx2A>trace 10.1.x1.49

Type escape sequence to abort.
Tracing the route to 10.1.x1.49

  1 150.x.x2.18 12 msec 12 msec 16 msec
  2 192.168.x.66 [MPLS: Labels 31/35 Exp 0] 340 msec 297 msec 388 msec
  3 192.168.100.** [MPLS: Labels 33/35 Exp 0] 249 msec 296 msec 349 msec
  4 192.168.100.** [MPLS: Labels 31/35 Exp 0] 248 msec 297 msec 300 msec
  5 150.x.x1.18 [AS 65001] [MPLS: Label 35 Exp 0] 149 msec 144 msec 148 msec
  6 150.x.x1.17 [AS 65001] 156 msec *  88 msec
```

Note	The third and fourth router addresses will vary, based on the SP.

Answer Key

The correct answers and expected solutions for the activities that are described in this guide appear here.

Lab 2-1 Answer Key: Establishing the Service Provider IGP Routing Environment

When you complete this activity, your router will be similar to the results here, with differences that are specific to your SP. The PE routers need the EIGRP network 150.*x*.0.0 command only for testing. It can be removed after testing. The CE routers will need network 150.*x*.0.0 in a later lab, so you could add the network statement in this lab.

Task 1: Configure the Service Provider IP Interfaces

Configuration steps on PE*x*1:

```
PEx1(config)#interface loopback0
PEx1(config-if)#ip address 192.168.x.17 255.255.255.255
PEx1(config-if)#no shutdown
PEx1(config-if)#interface serial 0/0
PEx1(config-if)#encapsulation frame-relay
PEx1(config-if)#frame-relay lmi-type cisco
PEx1(config-if)#interface Serial0/0.101 point-to-point
PEx1(config-subif)#ip address 150.x.x1.18 255.255.255.240
PEx1(config-subif)#frame-relay interface-dlci 101
PEx1(config-subif)#interface Serial0/0.102 point-to-point
PEx1(config-subif)#ip address 150.x.x1.34 255.255.255.240
PEx1(config-subif)#frame-relay interface-dlci 101
PEx1(config-subif)#interface Serial0/0.111 point-to-point
PEx1(config-subif)#ip address 192.168.x.49 255.255.255.240
PEx1(config-subif)#frame-relay interface-dlci 111
```

Configuration steps on PE*x*2:

```
PEx2(config)#interface loopback0
PEx2(config-if)#ip address 192.168.x.33 255.255.255.255
PEx2(config-if)#no shutdown
PEx2(config-if)#interface serial 0/0
PEx2(config-if)#encapsulation frame-relay
PEx2(config-if)#frame-relay lmi-type cisco
PEx2(config-if)#interface Serial0/0.101 point-to-point
PEx2(config-subif)#ip address 150.x.x2.18 255.255.255.240
PEx2(config-subif)#frame-relay interface-dlci 101
PEx2(config-subif)#interface Serial0/0.102 point-to-point
PEx2(config-subif)#ip address 150.x.x2.34 255.255.255.240
PEx2(config-subif)#frame-relay interface-dlci 101
PEx2(config-subif)#interface Serial0/0.111 point-to-point
PEx2(config-subif)#ip address 192.168.x.65 255.255.255.240
PEx2(config-subif)#frame-relay interface-dlci 111
```

Configuration steps on P*x*1:

```
Px1(config)#interface loopback0
Px1(config-if)#ip address 192.168.x.81 255.255.255.255
Px1(config-if)#no shutdown
Px1(config-if)#interface serial 0/0
Px1(config-if)#encapsulation frame-relay
Px1(config-if)#frame-relay lmi-type cisco
Px1(config-if)#interface Serial0/0.111 point-to-point
Px1(config-subif)#ip address 192.168.x.50 255.255.255.240
Px1(config-subif)#frame-relay interface-dlci 111
Px1(config-if)#interface Serial0/0.112 point-to-point
Px1(config-subif)#ip address 192.168.x.113 255.255.255.240
Px1(config-subif)#frame-relay interface-dlci 112
```

Configuration steps on P*x*2:

```
Px2(config)#interface loopback0
Px2(config-if)#ip address 192.168.x.97 255.255.255.255
Px2(config-if)#no shutdown
Px2(config-if)#interface serial 0/0
Px2(config-if)#encapsulation frame-relay
Px2(config-if)#frame-relay lmi-type cisco
Px2(config-if)#interface Serial0/0.111 point-to-point
Px2(config-subif)#ip address 192.168.x.66 255.255.255.240
Px2(config-subif)#frame-relay interface-dlci 111
Px2(config-if)#interface Serial0/0.112 point-to-point
Px2(config-subif)#ip address 192.168.x.114 255.255.255.240
Px2(config-subif)#frame-relay interface-dlci 112
```

Configuration steps on CE*x*1A routers:

```
CEx1A(config)#interface loopback0
CEx1A(config-if)#ip address 10.1.x1.49 255.255.255.255
CEx1A(config-if)#interface ethernet0/0
CEx1A(config-if)#ip address 10.1.x1.17 255.255.255.240
CEx1A(config-if)#no shutdown
CEx1A(config-if)#interface serial 0/0
CEx1A(config-if)#encapsulation frame-relay
CEx1A(config-if)#frame-relay lmi-type cisco
CEx1A(config-if)#interface Serial0/0.101 point-to-point
CEx1A(config-subif)#ip address 150.x.x1.17 255.255.255.240
CEx1A(config-subif)#frame-relay interface-dlci 101
```

Configuration steps on CE*x*1B routers:

```
CEx1B(config)#interface loopback0
CEx1B(config-if)#ip address 10.2.x1.49 255.255.255.255
CEx1B(config-if)#interface ethernet0/0
CEx1B(config-if)#ip address 10.2.x1.17 255.255.255.240
CEx1B(config-if)#no shutdown
CEx1B(config-if)#interface serial 0/0
CEx1B(config-if)#encapsulation frame-relay
CEx1B(config-if)#frame-relay lmi-type cisco
CEx1B(config-if)#interface Serial0/0.102 point-to-point
CEx1B(config-subif)#ip address 150.x.x1.33 255.255.255.240
CEx1B(config-subif)#frame-relay interface-dlci 102
```

Configuration steps on CE*x*2A routers:

```
CEx2A(config)#interface loopback0
CEx2A(config-if)#ip address 10.1.x2.49 255.255.255.255
CEx2A(config-if)#interface ethernet0/0
CEx2A(config-if)#ip address 10.1.x2.17 255.255.255.240
CEx2A(config-if)#no shutdown
CEx2A(config-if)#interface serial 0/0
CEx2A(config-if)#encapsulation frame-relay
CEx2A(config-if)#frame-relay lmi-type cisco
CEx2A(config-if)#interface Serial0/0.101 point-to-point
CEx2A(config-subif)#ip address 150.x.x2.17 255.255.255.240
CEx2A(config-subif)#frame-relay interface-dlci 101
```

Configuration steps on CEx2B routers:

```
CEx2B(config)#interface loopback0
CEx2B(config-if)#ip address 10.2.x2.49 255.255.255.255
CEx2B(config-if)#interface ethernet0/0
CEx2B(config-if)#ip address 10.2.x2.17 255.255.255.240
CEx2B(config-if)#no shutdown
CEx2B(config-if)#interface serial 0/0
CEx2B(config-if)#encapsulation frame-relay
CEx2B(config-if)#frame-relay lmi-type cisco
CEx2B(config-if)#interface Serial0/0.102 point-to-point
CEx2B(config-subif)#ip address 150.x.x2.33 255.255.255.240
CEx2B(config-subif)#frame-relay interface-dlci 102
```

Task 2: Configure the Service Provider IGP

Configuration steps on PE*x*1:

```
PEx1(config)#router eigrp 1
PEx1(config-router)#network 150.x.0.0 (optional)
PEx1(config-router)#network 192.168.x.0
PEx1(config-router)#no auto-summary
```

Configuration steps on PE*x*2:

```
PEx2(config)#router eigrp 1
PEx2(config-router)#network 150.x.0.0 (optional)
PEx2(config-router)#network 192.168.x.0
PEx2(config-router)#no auto-summary
```

Configuration steps on P*x*1:

```
Px1(config)#router eigrp 1
Px1(config-router)#network 192.168.x.0
Px1(config-router)#no auto-summary
```

Configuration steps on P*x*2:

```
Px2(config)#router eigrp 1
Px2(config-router)#network 192.168.x.0
Px2(config-router)#no auto-summary
```

Configuration steps on all CE routers:

```
CEx**(config)#router rip
CEx**(config-router)#network 10.0.0.0
CEx**(config-router)#network 150.x.0.0 (optional)
CEx**(config-router)#no auto-summary
```

Lab 3-1 Answer Key: Establishing the Core MPLS Environment

When you complete this activity, your router will be similar to the results here, with differences that are specific to your SP.

Task 1: Enable LDP on Your PE and P Routers

Configuration steps on PE*x*1:

```
PEx1(config)#ip cef
PEx1(config)#interface serial0/0.111
PEx1(config-subif)#mpls label protocol ldp
PEx1(config-subif)#mpls ip
```

Note The **mpls label protocol ldp** command can be issued at the global configuration level.

Configuration steps on PE*x*2:

```
PEx2(config)#ip cef
PEx2(config)#interface serial0/0.111
PEx2(config-subif)#mpls label protocol ldp
PEx2(config-subif)#mpls ip
```

Configuration steps on P*x*1:

```
Px1(config)#ip cef
Px1(config)#interface serial0/0.111
Px1(config-subif)#mpls label protocol ldp
Px1(config-subif)#mpls ip
Px1(config)#interface serial0/0.112
Px1(config-subif)#mpls label protocol ldp
Px1(config-subif)#mpls ip
```

Configuration steps on P*x*2:

```
Px2(config)#ip cef
Px2(config)#interface serial0/0.111
Px2(config-subif)#mpls label protocol ldp
Px2(config-subif)#mpls ip
Px2(config)#interface serial0/0.112
Px2(config-subif)#mpls label protocol ldp
Px2(config-subif)#mpls ip
```

Note The **mpls ip** command is issued to enable MPLS on an interface but will be displayed in the **show running-config** command output as the **tag-switching ip** command.

Task 2: Experiment with TTL Propagation

Configuration steps on PE*x*1 and PEx2 to enable MPLS TTL propagation:

```
PEx*(config)#mpls ip propagate-ttl
```

Configuration steps on P*x*1 and P*x*2 to enable MPLS TTL propagation:

```
Px*(config)#mpls ip propagate-ttl
```

Configuration steps on PE*x*1 and PE*x*2 to disable MPLS TTL propagation:

```
PEx*(config)#no mpls ip propagate-ttl
```

Configuration steps on P*x*1 and P*x*2 to disable MPLS TTL propagation:

```
Px*(config)#no mpls ip propagate-ttl
```

Task 3: Configure Conditional Label Distribution

Note There are various ways to construct an access list to accomplish the desired result. This example shows one way. The key, however, is to meet the task objective.

Configuration steps on PE*x*1 and PE*x*2:

```
PEx*(config)#no mpls ldp advertise-labels
PEx*(config)#mpls ldp advertise-labels for 90
PEx*(config)#access-list 90 permit 150.x.0.0 0.0.255.255
PEx*(config)#access-list 90 permit 192.168.x.16 0.0.0.15
PEx*config)#access-list 90 permit 192.168.x.32 0.0.0.15
PEx*(config)#access-list 90 permit 192.168.x.80 0.0.0.15
PEx*(config)#access-list 90 permit 192.168.x.96 0.0.0.15
```

Configuration steps on P*x*1 and P*x*2:

```
Px*(config)#no mpls ldp advertise-labels
Px*(config)#mpls ldp advertise-labels for 90
Px*(config)#access-list 90 permit 150.x.0.0 0.0.255.255
Px*(config)#access-list 90 permit 192.168.x.16 0.0.0.15
Px*(config)#access-list 90 permit 192.168.x.32 0.0.0.15
Px*(config)#access-list 90 permit 192.168.x.80 0.0.0.15
Px*(config)#access-list 90 permit 192.168.x.96 0.0.0.15
```

Task 4: Remove Conditional Label Distribution

Configuration steps on PE*x*1 and PE*x*2:

```
PEx*(config)#no mpls ldp advertise-labels for 90
PEx*(config)#mpls ldp advertise-labels
PEx*(config)#no access-list 90
```

Configuration steps on P*x*1 and P*x*2:

```
Px*(config)#no mpls ldp advertise-labels for 90
Px*(config)#mpls ldp advertise-labels
Px*(config)#no access-list 90
```

Lab 5-1 Answer Key: Configuring Initial MPLS VPN Setup

When you complete this activity, your router will be similar to the results here, with differences that are specific to your SP.

Task 1: Configure MP-BGP

Configuration steps on PE*x*1:

```
PEx1(config)#router bgp 65001
PEx1(config-router)#neighbor 192.168.x.33 remote-as 65001
PEx1(config-router)#neighbor 192.168.x.33 update-source loopback 0
PEx1(config-router)#no auto-summary
PEx1(config-router)#address-family vpnv4
PEx1(config-router-af)#neighbor 192.168.x.33 activate
PEx1(config-router-af)#neighbor 192.168.x.33 next-hop-self
PEx1(config-router-af)#neighbor 192.168.x.33 send-community both
PEx1(config-router-af)#no auto-summary
```

Configuration steps on PE*x*2:

```
PEx2(config)#router bgp 65001
PEx2(config-router)#neighbor 192.168.x.17 remote-as 65001
PEx2(config-router)#neighbor 192.168.x.17 update-source loopback 0
PEx2(config-router)#no auto-summary
PEx2(config-router)#address-family vpnv4
PEx2(config-router-af)#neighbor 192.168.x.17 activate
PEx2(config-router-af)#neighbor 192.168.x.17 next-hop-self
PEx2(config-router-af)#neighbor 192.168.x.17 send-community both
PEx2(config-router-af)#no auto-summary
```

Task 2: Configure VRF Tables

Configuration steps on PE*x*1:

```
PEx1(config)#ip vrf CustA
PEx1(config-vrf)#rd x:10
PEx1(config-vrf)#route-target both x:10
PEx1(config)#ip vrf CustB
PEx1(config-vrf)#rd x:20
PEx1(config-vrf)#route-target both x:20
PEx1(config)#interface serial0/0.101
PEx1(config-subif)#ip vrf forwarding CustA
PEx1(config-subif)#ip address 150.x.x1.18 255.255.255.240
PEx1(config)#int serial0/0.102
PEx1(config-subif)#ip vrf forwarding CustB
PEx1(config-subif)#ip address 150.x.x1.34 255.255.255.240
PEx1(config)#router rip
PEx1(config-router)#version 2
```

```
PEx1(config-router)#address-family ipv4 vrf CustA
PEx1(config-router-af)#network 150.x.0.0
PEx1(config-router-af)#no auto-summary
PEx1(config-router-af)#redistribute bgp 65001 metric transparent
PEx1(config-router)#address-family ipv4 vrf CustB
PEx1(config-router-af)#network 150.x.0.0
PEx1(config-router-af)#no auto-summary
PEx1(config-router-af)#redistribute bgp 65001 metric transparent
PEx1(config-router)#router bgp 65001
PEx1(config-router)#address-family ipv4 vrf CustA
PEx1(config-router-af)#no auto-summary
PEx1(config-router-af)#redistribute rip
PEx1(config-router-af)#exit
PEx1(config-router)#address-family ipv4 vrf CustB
PEx1(config-router-af)#no auto-summary
PEx1(config-router-af)#redistribute rip
```

Configuration steps on PE*x*2:

```
PEx2(config)#ip vrf CustA
PEx2(config-vrf)#rd x:10
PEx2(config-vrf)#route-target both x:10
PEx2(config)#ip vrf CustB
PEx2(config-vrf)#rd x:20
PEx2(config-vrf)#route-target both x:20
PEx2(config)#interface serial0/0.101
PEx2(config-subif)#ip vrf forwarding CustA
PEx2(config-subif)#ip address 150.x.x2.18 255.255.255.240
PEx2(config)#interface serial0/0.102
PEx2(config-subif)#ip vrf forwarding CustB
PEx2(config-subif)#ip address 150.x.x2.34 255.255.255.240
PEx2(config)#router rip
PEx2(config-router)#version 2
PEx2(config-router)#address-family ipv4 vrf CustA
PEx2(config-router-af)#network 150.x.0.0
PEx2(config-router-af)#no auto-summary
PEx2(config-router-af)#redistribute bgp 65001 metric transparent
PEx2(config-router)#address-family ipv4 vrf CustB
PEx2(config-router-af)#network 150.x.0.0
PEx2(config-router-af)#no auto-summary
PEx2(config-router-af)#redistribute bgp 65001 metric transparent
PEx2(config)#router bgp 65001
PEx2(config-router)#address-family ipv4 vrf CustA
PEx2(config-router-af)#no auto-summary
PEx2(config-router-af)#redistribute rip
PEx2(config-router)#address-family ipv4 vrf CustB
PEx2(config-router-af)#no auto-summary
PEx2(config-router-af)#redistribute rip
```

Lab 5-2 Answer Key: Running EIGRP Between PE and CE Routers

When you complete this activity, your router will be similar to the results here, with differences that are specific to your SP.

Task 1: Enable an EIGRP VPN

Configuration steps on CE*x*1A:

```
CEx1A(config)#no router rip
CEx1A(config)#router eigrp x
CEx1A(config-router)#network 10.0.0.0
CEx1A(config-router)#network 150.x.0.0
CEx1A(config-router)#no auto-summary
```

Configuration steps on CE*x*2B:

```
CEx2B(config)#no router rip
CEx2B(config)#router eigrp x
CEx2B(config-router)#network 10.0.0.0
CEx2B(config-router)#network 150.x.0.0
CEx2B(config-router)#no auto-summary
```

Configuration steps on PE*x*1:

```
PEx1(config)#router rip
PEx1(config-router)#no address-family ipv4 vrf CustA
PEx1(config)#router eigrp 1
PEx1(config-router)#address-family ipv4 vrf CustA
PEx1(config-router-af)#autonomous-system x
PEx1(config-router-af)#network 150.x.x1.16 0.0.0.15
PEx1(config-router-af)#no auto-summary
PEx1(config-router-af)#redistribute bgp 65001 metric 10000 100 255 1 1500
PEx1(config-router-af)#exit
PEx1(config-router)#router bgp 65001
PEx1(config-router)#address-family ipv4 vrf CustA
PEx1(config-router-af)#no redistribute rip
PEx1(config-router-af)#redistribute eigrp x metric 1
```

Configuration steps on PE*x*2:

```
PEx2(config)#router rip
PEx2(config-router)#no address-family ipv4 vrf CustB
PEx2(config-router)#router eigrp 1
PEx2(config-router)#address-family ipv4 vrf CustB
PEx2(config-router-af)#autonomous-system x
PEx2(config-router-af)#network 150.x.x2.32 0.0.0.15
PEx2(config-router-af)#no auto-summary
PEx2(config-router-af)#redistribute bgp 65001 metric 10000 100 255 1 1500
```

```
PEx2(config-router-af)#exit
PEx2(config-router)#router bgp 65001
PEx2(config-router)#address-family ipv4 vrf CustB
PEx2(config-router-af)#no redistribute rip
PEx2(config-router-af)#redistribute eigrp x metric 1
```

Lab 5-3 Answer Key: Running OSPF Between PE and CE Routers

When you complete this activity, your router will be similar to the results here, with differences that are specific to your SP.

Task 1: Configure OSPF as the PE-CE Routing Protocol

Configuration steps on CE*x*1A:

```
CEx1A(config)#no router eigrp x
CEx1A(config)#router ospf 1
CEx1A(config-router)#network 150.x.0.0 0.0.255.255 area 0
CEx1A(config-router)#network 10.1.x1.49 0.0.0.0 area 0
CEx1A(config-router)#network 10.1.x1.16 0.0.0.15 area 1
```

Configuration steps on CE*x*1B:

```
CEx1B(config)#no router rip
CEx1B(config)#router ospf 2
CEx1B(config-router)#network 150.x.0.0 0.0.255.255 area 0
CEx1B(config-router)#network 10.2.x1.49 0.0.0.0 area 0
CEx1B(config-router)#network 10.2.x1.16 0.0.0.15 area 1
```

Configuration steps on CE*x*2A:

```
CEx2A(config)#no router rip
CEx2A(config)#router ospf 1
CEx2A(config-router)#network 150.x.0.0 0.0.255.255 area 0
CEx2A(config-router)#network 10.1.x2.49 0.0.0.0 area 0
CEx2A(config-router)#network 10.1.x2.16 0.0.0.15 area 1
```

Configuration steps on CE*x*2B:

```
CEx2B(config)#no router eigrp x
CEx2B(config)#router ospf 2
CEx2B(config-router)#network 150.x.0.0 0.0.255.255 area 0
CEx2B(config-router)#network 10.2.x2.49 0.0.0.0 area 0
CEx2B(config-router)#network 10.2.x2.16 0.0.0.15 area 1
```

Configuration steps on PE*x*1:

```
PEx1(config)#no router rip
PEx1(config)#router ospf 2 vrf CustB
PEx1(config-router)#network 150.x.0.0 0.0.255.255 area 0
PEx1(config-router)#redistribute bgp 65001 subnets
PEx1(config-router)#exit
PEx1(config)#router bgp 65001
PEx1(config-router)#address-family ipv4 vrf CustB
PEx1(config-router)#no redistribute rip
PEx1(config-router-af)#redistribute ospf 2
PEx1(config-router-af)#exit
```

```
PEx1(config)#router eigrp 1
PEx1(config-router)#no address-family ipv4 vrf CustA
PEx1(config)#router ospf 1 vrf CustA
PEx1(config-router)#network 150.x.0.0 0.0.255.255 area 0
PEx1(config-router)#redistribute bgp 65001 subnets
PEx1(config-router)#exit
PEx1(config)#router bgp 65001
PEx1(config-router)#address-family ipv4 vrf CustA
PEx1(config-router-af)#no redistribute eigrp x
PEx1(config-router-af)#redistribute ospf 1
```

Configuration steps on PE*x*2:

```
PEx2(config)#no router rip
PEx2(config)#router ospf 1 vrf CustA
PEx2(config-router)#network 150.x.0.0 0.0.255.255 area 0
PEx2(config-router)#redistribute bgp 65001 subnets
PEx2(config-router)#exit
PEx2(config)#router bgp 65001
PEx2(config-router)#address-family ipv4 vrf CustA
PEx2(config-router)#no redistribute rip
PEx2(config-router-af)#redistribute ospf 1
PEx2(config-router-af)#exit
PEx2(config)#router eigrp 1
PEx2(config-router)#no address-family ipv4 vrf CustB
PEx2(config)#router ospf 2 vrf CustB
PEx2(config-router)#network 150.x.0.0 0.0.255.255 area 0
PEx2(config-router)#redistribute bgp 65001 subnets
PEx2(config-router)#exit
PEx2(config)#router bgp 65001
PEx2(config-router)#address-family ipv4 vrf CustB
PEx2(config-router-af)#no redistribute eigrp x
PEx2(config-router-af)#redistribute ospf 2
```

Lab 5-4 Answer Key: Running BGP Between PE and CE Routers

When you complete this activity, your router will be similar to the results here, with differences that are specific to your SP.

Task 1: Configure BGP as the PE-CE Routing Protocol

Configuration steps on CE*x*1A:

```
CEx1A(config)#router bgp 650x1
CEx1A(config-router)#neighbor 150.x.x1.18 remote-as 65001
CEx1A(config-router)#no auto-summary
CEx1A(config-router)#redistribute ospf 1
CEx1A(config)#router ospf 1
CEx1A(config-router)#redistribute bgp 650x1 subnets
```

Configuration steps on CE*x*1B:

```
CEx1B(config)#router bgp 650x2
CEx1B(config-router)#neighbor 150.x.x1.34 remote-as 65001
CEx1B(config-router)#no auto-summary
CEx1B(config-router)#redistribute ospf 2
CEx1B(config-router)#router ospf 2
CEx1B(config-router)#redistribute bgp 650x2 subnets
```

Configuration steps on CE*x*2A:

```
CEx2A(config)#router bgp 650x1
CEx2A(config-router)#neighbor 150.x.x2.18 remote-as 65001
CEx2A(config-router)#no auto-summary
CEx2A(config-router)#redistribute ospf 1
CEx2A(config-router)#router ospf 1
CEx2A(config-router)#redistribute bgp 650x1 subnets
```

Configuration steps on CE*x*2B:

```
CEx2B(config)#router bgp 650x2
CEx2B(config-router)#neighbor 150.x.x2.34 remote-as 65001
CEx2B(config-router)#no auto-summary
CEx2B(config-router)#redistribute ospf 2
CEx2B(config-router)#router ospf 2
CEx2B(config-router)#redistribute bgp 650x2 subnets
```

Configuration steps on PE*x*1:

```
!******* PoP 1 *************
PEx1(config)#router bgp 65001
PEx1(config-router)#address-family ipv4 vrf CustA
PEx1(config-router-af)#no redistribute ospf 1
PEx1(config)#no router ospf 1 vrf CustA
PEx1(config)#router bgp 65001
PEx1(config-router)#address-family ipv4 vrf CustA
PEx1(config-router-af)#neighbor 150.x.x1.17 remote-as 650x1
PEx1(config-router-af)#neighbor 150.x.x1.17 activate
PEx1(config-router-af)#neighbor 150.x.x1.17 as-override

!******* PoP 2 **************
PEx1(config)#router bgp 65001
PEx1(config-router-af)#address-family ipv4 vrf CustB
PEx1(config-router-af)#no redistribute ospf 2
PEx1(config)#no router ospf 2 vrf CustB
PEx1(config)#router bgp 65001
PEx1(config-router-af)#address-family ipv4 vrf CustB
PEx1(config-router-af)#neighbor 150.x.x1.33 remote-as 650x2
PEx1(config-router-af)#neighbor 150.x.x1.33 activate
PEx1(config-router-af)#neighbor 150.x.x1.33 as-override
```

Configuration steps on PE*x*2:

```
!******* PoP 1 *******
PEx2(config)#router bgp 65001
PEx2(config-router)#address-family ipv4 vrf CustA
PEx2(config-router-af)#no redistribute ospf 1
PEx2(config)#no router ospf 1 vrf CustA
PEx2(config)#router bgp 65001
PEx2(config-router)#address-family ipv4 vrf CustA
PEx2(config-router-af)#neighbor 150.x.x2.17 remote-as 650x1
PEx2(config-router-af)#neighbor 150.x.x2.17 activate
PEx2(config-router-af)#neighbor 150.x.x2.17 as-override

!******* PoP 2 *******
PEx2(config-router-af)#address-family ipv4 vrf CustB
PEx2(config-router-af)#no redistribute ospf 2
PEx2(config)#no router ospf 2 vrf CustB
PEx2(config)#router bgp 65001
PEx2(config-router-af)#address-family ipv4 vrf CustB
PEx2(config-router-af)#neighbor 150.x.x2.33 remote-as 650x2
PEx2(config-router-af)#neighbor 150.x.x2.33 activate
PEx2(config-router-af)#neighbor 150.x.x2.33 as-override
```

Task 2: Configure the Backup PE-CE Link

Answer to "Which VRF is interface Se0/0.113 from CE*x*1B added to?" CustB

Answer to "Which VRF is interface Se0/0.113 from CE*x*2A added to?" CustA

Configuration steps on CE*x*1B:

```
CEx1B(config)#interface serial0/0.113 point-to-point
CEx1B(config-subif)#ip address 150.x.x2.49 255.255.255.240
CEx1B(config-subif)#frame-relay interface-dlci 113
CEx1B(config-fr-dlci)#no shut
CEx1B(config)#router bgp 650x2
CEx1B(config-router)#neighbor 150.x.x2.50 remote-AS 65001
```

Configuration steps on PE*x*2:

```
PEx2(config)#interface serial0/0.113 point-to-point
PEx2(config-subif)#ip vrf forwarding CustB
PEx2(config-subif)#ip address 150.x.x2.50 255.255.255.240
PEx2(config-subif)#frame-relay interface-dlci 113
PEx2(config-fr-dlci)#no shut
PEx2(config)#router bgp 65001
PEx2(config-router-af)#address-family ipv4 vrf CustB
PEx2(config-router-af)#neighbor 150.x.x2.49 remote-as 650x2
PEx2(config-router-af)#neighbor 150.x.x2.49 activate
PEx2(config-router-af)#neighbor 150.x.x2.49 as-override
```

Configuration steps on CE*x*2A:

```
CEx2A(config)#interface serial0/0.113 point-to-point
CEx2A(config-subif)#ip address 150.x.x1.49 255.255.255.240
CEx2A(config-subif)#frame-relay interface-dlci 113
CEx2A(config-fr-dlci)#no shut
CEx2A(config)#router bgp 650x1
CEx2A(config-router)#neighbor 150.x.x1.50 remote-as 65001
```

Configuration steps on PE*x*1:

```
PEx1(config)#interface serial0/0.113 point-to-point
PEx1(config-subif)#ip vrf forwarding CustA
PEx1(config-subif)#ip address 150.x.x1.50 255.255.255.240
PEx1(config-subif)#frame-relay interface-dlci 113
PEx1(config-fr-dlci)#no shut
PEx1(config)#router bgp 65001
PEx1(config-router)#address-family ipv4 vrf CustA
PEx1(config-router-af)#neighbor 150.x.x1.49 remote-as 650x1
PEx1(config-router-af)#neighbor 150.x.x1.49 activate
PEx1(config-router-af)#neighbor 150.x.x1.49 as-override
```

Task 3: Select the Primary and Backup Link with BGP

Configuration steps on CE*x*1B:

```
CEx1B(config)#route-map setLP permit 10
CEx1B(config-route-map)#set local-preference 50
CEx1B(config-route-map)#route-map setMED permit 10
CEx1B(config-route-map)#set metric 200
CEx1B(config-route-map)#router bgp 650x2
CEx1B(config-router)#neighbor 150.x.x2.50 route-map setLP in
CEx1B(config-router)#neighbor 150.x.x2.50 route-map setMED out
```

Configuration steps on CE*x*2A:

```
CEx2A(config)#route-map setLP permit 10
CEx2A(config-route-map)#set local-preference 50
CEx2A(config-route-map)#route-map setMED permit 10
CEx2A(config-route-map)#set metric 200
CEx2A(config-route-map)#router bgp 650x1
CEx2A(config-router)#neighbor 150.x.x1.50 route-map setLP in
CEx2A(config-router)#neighbor 150.x.x1.50 route-map setMED out
```

Lab 6-1 Answer Key: Establishing Overlapping VPNs

When you complete this activity, your router will be similar to the results here, with differences that are specific to your SP.

Task 1: Design Your VPN Solution

No configuration steps are required for this task. On paper, you defined new RDs for VRFs to which CE*x*1A and CE*x*2B will be connected. A new RT is needed for the CustAB VPN. The RT value should be the same with the other POP within your SP. It was suggested to use *x*:11 as the RD for VRFs connected to CE*x*1A, *x*:21 as the RD for VRFs connected to CE*x*2B, and *x*:1001 as the RT for the CustAB VPN.

Task 2: Remove CE*x*1A and CE*x*2B from Existing VRFs

Configuration steps on PE*x*1:

```
PEx1(config)#router bgp 65001
PEx1(config-router)#address-family ipv4 vrf CustA
PEx1(config-router-af)#no neighbor 150.x.x1.17
PEx1(config-vrf)#interface serial0/0.101
PEx1(config-subif)#no ip vrf forwarding CustA
```

Note	After removing the interface from the VRF, the following message will appear: "% Interface Serial0/0.101 IP address 150.*x*.*x*1.18 removed due to disabling VRF CustA."

Configuration steps on PE*x*2:

```
PEx2(config)#router bgp 65001
PEx2(config-router)#address-family ipv4 vrf CustB
PEx2(config-router-af)#no neighbor 150.x.x2.33
PEx2(config-vrf)#interface serial0/0.102
PEx2(config-subif)#no ip vrf forwarding CustB
```

Note	After removing the interface from the VRF, the following message will appear: "% Interface Serial0/0.102 IP address 150.*x*.*x*2.34 removed due to disabling VRF CustB."

Task 3: Configure New VRFs for CE*x*1A and CE*x*2B

Note	The RDs and RTs listed in these results may not match those that you have used in this task.

Configuration steps on PE*x*1:

```
PEx1(config)#ip vrf CustAB
PEx1(config-vrf)#rd x:11
PEx1(config-vrf)#route-target both x:10
```

```
PEx1(config-vrf)#route-target both x:1001
PEx1(config-vrf)#interface serial0/0.101
PEx1(config-subif)#ip vrf forwarding CustAB
PEx1(config-subif)#ip address 150.x.x1.18 255.255.255.240
PEx1(config)#router bgp 65001
PEx1(config-router-af)#address-family ipv4 vrf CustAB
PEx1(config-router-af)#neighbor 150.x.x1.17 remote-as 650x1
PEx1(config-router-af)#neighbor 150.x.x1.17 activate
PEx1(config-router-af)#neighbor 150.x.x1.17 as-override
```

Configuration steps on PE*x2*:

```
PEx2(config)#ip vrf CustAB
PEx2(config-vrf)#rd x:21
PEx2(config-vrf)#route-target both x:20
PEx2(config-vrf)#route-target both x:1001
PEx2(config-vrf)#interface serial0/0.102
PEx2(config-subif)#ip vrf forwarding CustAB
PEx2(config-subif)#ip address 150.x.x2.34 255.255.255.240
PEx2(config)#router bgp 65001
PEx2(config-router-af)#address-family ipv4 vrf CustAB
PEx2(config-router-af)#neighbor 150.x.x2.33 remote-as 650x2
PEx2(config-router-af)#neighbor 150.x.x2.33 activate
PEx2(config-router-af)#neighbor 150.x.x2.33 as-override
```

Lab 6-2 Answer Key: Merging Service Providers

When you complete this activity, your router will be similar to the results here, with differences that are specific to your SP.

Task 1: Enable Connectivity with the Central P Router

Note	The subinterface number and DLCI number in these configurations will match each other and are determined by the instructions for this task. This task is for the P router of the SP, and not the PE router.

Configuration steps on P*x*1:

```
Px1(config)#interface serial0/0.2x1 point-to-point
Px1(config-subif)#ip address 192.168.100.** 255.255.255.248
Px1(config-subif)#frame-relay interface-dlci 2x1
Px1(config-fr-dlci)#no shut
```

Configuration steps on P*x*2:

```
Px2(config)#interface serial0/0.2x2 point-to-point
Px2(config-subif)#ip address 192.168.100.** 255.255.255.248
Px2(config-subif)#frame-relay interface-dlci 2x2
Px2(config-fr-dlci)#no shut
```

Task 2: Migrate the Core to IS-IS

Note	Take care to apply the configuration steps to the appropriate PE or P router.

Configuration steps on PE*x*1:

```
PEx1(config)#no router eigrp 1
PEx1(config)#router isis
PEx1(config-router)#net 49.0001.0000.0000.01x1.00
PEx1(config-router)#is-type level-2-only
PEx1(config-router)#metric-style wide
PEx1(config-router)#interface serial0/0.111
PEx1(config-subif)#ip router isis
PEx1(config)#interface loopback0
PEx1(config-subif)#ip router isis
```

Configuration steps on PE*x*2:

```
Px2(config)#no router eigrp 1
PEx2(config)#router isis
PEx2(config-router)#net 49.0001.0000.0000.01x2.00
PEx2(config-router)#is-type level-2-only
PEx2(config-router)#metric-style wide
PEx2(config)#interface serial0/0.111
PEx2(config-subif)#ip router isis
PEx2(config)#interface loopback0
PEx2(config-subif)#ip router isis
```

Note	The configuration steps shown next apply to the P routers.

Configuration steps on P*x*1:

```
Px1(config)#no router eigrp 1
Px1(config)#router isis
Px1(config-router)#net 49.0001.0000.0000.02x1.00
Px1(config-router)#is-type level-2-only
Px1(config-router)#metric-style wide
Px1(config-router)#interface serial0/0.111
Px1(config-subif)#ip router isis
Px1(config-router)#interface serial0/0.112
Px1(config-subif)#ip router isis
Px1(config-router)#interface serial0/0.2x1
Px1(config-subif)#ip router isis
Px1(config)#interface loopback0
Px1(config-subif)#ip router isis
```

Configuration steps on P*x*2:

```
Px2(config)#no router eigrp 1
Px2(config)#router isis
Px2(config-router)#net 49.0001.0000.0000.02x2.00
Px2(config-router)#is-type level-2-only
Px2(config-router)#metric-style wide
Px2(config)#interface serial0/0.111
Px2(config-subif)#ip router isis
Px2(config-router)#interface serial0/0.112
Px2(config-subif)#ip router isis
Px2(config-router)#interface serial0/0.2x2
Px2(config-subif)#ip router isis
Px2(config)#interface loopback0
Px2(config-subif)#ip router isis
```

Task 3: Enable MPLS LDP Connectivity with the Central P Router

Note	The subinterface number and DLCI number in these configurations match each other and are determined by the instructions for this task.

Configuration steps on P*x*1:

```
Px1(config)#interface serial0/0.2x1
Px1(config-subif)#mpls ip
Px1(config-subif)#mpls label protocol ldp
```

Configuration steps on P*x*2:

```
Px2(config)#interface serial0/0.2x2
Px2(config-subif)#mpls ip
Px2(config-subif)#mpls label protocol ldp
```

Task 4: Enable IBGP Connectivity for All PE Routers

Answer to "What routes do you expect to see in VRF CustA?"

Only routes from your SP: 10.1.x1.16/28, 10.1.x1.49/32, 10.1.x2.16/28, 10.1.x2.49/32, 150.x.x1.16/18, 150.x.x1.48/28, 150.x.x2.16/18

Answer to "What routes do you expect to see in VRF CustB?"

Only routes from your SP: 10.2.x1.16/28, 10.2.x1.49/32, 10.2.x2.16/28, 10.2.x2.49/32, 150.x.x2.32/18, 150.x.x2.48/28, 150.x.x1.32/18

Answer to "What routes do you expect to see in VRF CustAB?"

Only routes from your SP: 10.1.x1.16/28, 10.1.x1.49/32, 10.1.x2.16/28, 10.1.x2.49/32, 10.2.x2.16/28, 10.2.x2.49/32, 150.x.x1.16/18, 150.x.x1.48/28, 150.x.x2.16/18, , 150.x.x2.32/18

Configuration steps on PE*x*1:

```
PEx1(config)#router bgp 65001
PEx1(config-router)#no neighbor 192.168.x.33 remote-as 65001
PEx1(config-router)#neighbor 192.168.100.129 remote-as 65001
PEx1(config-router)#neighbor 192.168.100.129 update-source loopback0
PEx1(config-router)#address-family vpnv4
PEx1(config-router-af)#neighbor 192.168.100.129 activate
PEx1(config-router-af)#neighbor 192.168.100.129 send-community both
PEx1(config-router-af)#neighbor 192.168.100.129  next-hop-self
```

Configuration steps on PE*x*2:

```
PEx2(config)#router bgp 65001
PEx2(config-router)#no neighbor 192.168.x.17 remote-as 65001
PEx2(config-router)#neighbor 192.168.100.129 remote-as 65001
PEx2(config-router)#neighbor 192.168.100.129 update-source loopback0
PEx2(config-router)#address-family vpnv4
PEx2(config-router-af)#neighbor 192.168.100.129 act
PEx2(config-router-af)#neighbor 192.168.100.129 send-community both
PEx2(config-router-af)#neighbor 192.168.100.129  next-hop-self
```

Lab 6-3 Answer Key: Establishing a Common Services VPN

When you complete this activity, your router will be similar to the results here, with differences that are specific to your SP.

Task 1: Establish Connectivity Between the NMS VRF and Other VRFs

Answer to "What routes do you expect to see in VRF CustA on your PE?"

All the routes from before, plus the prefix for the new link and 10.10.10.49. These should include: 10.1.*x*1.16/28, 10.1.*x*1.49/32, 10.1.*x*2.16/28, 10.1.*x*2.49/32, 150.*x*.*x*1.16/18, 150.*x*.*x*1.48/28, 150.*x*.*x*1.64/28, 150.*x*.*x*2.16/18, and 10.10.10.10.49/32.

Answer to "What routes do you expect to see in VRF NMS on your PE?"

None, because the PE router has no interfaces in the NMS VRF and does not have the NMS VRF configured.

Answer to "What routes do you expect to see in VRF NMS on P1?"

The loopback addresses of the CE routers in the SP and 10.10.10.49. You need only check for your loopback address: 10.1.*x*1.49/32, 10.2.*x*1.49/32, 10.1.*x*2.49/32, 10.2.*x*1.49/32, and 10.10.10.49/32.

Configuration steps on PE*x*1 for customer A:

```
PEx1(config)#ip vrf CustA
PEx1(config-vrf)#export map NMS_Cus_A
PEx1(config-vrf)#route-target import 101:500
PEx1(config)#ip vrf CustAB
PEx1(config-vrf)#export map NMS_Cus_A
PEx1(config-vrf)#route-target import 101:500
PEx1(config)#route-map NMS_Cus_A permit 10
PEx1(config-route-map)#match ip address 10
PEx1(config-route-map)#set extcommunity rt 101:501 add
PEx1(config-route-map)#exit
PEx1(config)#access-list 10 permit host 10.1.x1.49
PEx1(config)#access-list 10 permit host 10.1.x2.49
```

Configuration steps on PE*x*2 for customer A:

```
PEx2(config)#ip vrf CustA
PEx2(config-vrf)#export map NMS_Cus_A
PEx2(config-vrf)#route-target import 101:500
PEx2(config)#route-map NMS_Cus_A permit 10
PEx2(config-route-map)#match ip address 10
PEx2(config-route-map)#set extcommunity rt 101:501 add
PEx2(config-route-map)#exit
```

```
PEx2(config)#access-list 10 permit host 10.1.x1.49
PEx2(config)#access-list 10 permit host 10.1.x2.49
```

Configuration steps on PE*x*1 for customer B:

```
PEx1(config)#ip vrf CustB
PEx1(config-vrf)#export map NMS_Cus_B
PEx1(config-vrf)#route-target import 101:500
PEx1(config)#route-map NMS_Cus_B permit 10
PEx1(config-route-map)#match ip address 20
PEx1(config-route-map)#set extcommunity rt 101:501 add
PEx1(config-route-map)#exit
PEx1(config)#access-list 20 permit host 10.2.x1.49
PEx1(config)#access-list 20 permit host 10.2.x2.49
```

Configuration steps on PE*x*2 for customer B:

```
PEx2(config)#ip vrf CustB
PEx2(config-vrf)#export map NMS_Cus_B
PEx2(config-vrf)#route-target import 101:500
PEx2(config)#ip vrf CustAB
PEx2(config-vrf)#export map NMS_Cus_B
PEx2(config-vrf)#route-target import 101:500
PEx2(config)#route-map NMS_Cus_B permit 10
PEx2(config-route-map)#match ip address 20
PEx2(config-route-map)#set extcommunity rt 101:501 add
PEx2(config-route-map)#exit
PEx2(config)#access-list 20 permit host 10.2.x1.49
PEx2(config)#access-list 20 permit host 10.2.x2.49
```

Lab 7-1 Answer Key: Establishing Central Site Internet Connectivity with an MPLS VPN

When you complete this activity, your router will be similar to the results here, with differences that are specific to your SP.

Task 1: Restore a Simple Customer VPN Configuration

Configuration steps on PE*x*1:

```
PEx1(config)#router bgp 65001
PEx1(config-router)#no address-family ipv4 vrf CustAB
PEx1(config-router)#!
PEx1(config-router)# address-family ipv4 vrf CustA
PEx1(config-router-af)# neighbor 150.x.x1.17 remote-as 650x1
PEx1(config-router-af)# neighbor 150.x.x1.17 activate
PEx1(config-router-af)# neighbor 150.x.x1.17 as-override
PEx1(config-router-af)#  exit-address-family
PEx1(config-router)# exit
PEx1(config)#no ip vrf CustAB
PEx1(config)#int ser 0/0.101
PEx1(config-subif)#ip vrf forwarding CustA
PEx1(config-subif)#ip add 150.x.x1.18 255.255.255.240
```

Configuration steps on PE*x*2:

```
PEx2(config)#router bgp 65001
PEx2(config-router)#no address-family ipv4 vrf CustAB
PEx2(config-router)#address-family ipv4 vrf CustB
PEx2(config-router-af)# neighbor 150.x.x2.33 remote-as 650x2
PEx2(config-router-af)# neighbor 150.x.x2.33 activate
PEx2(config-router-af)# neighbor 150.x.x2.33 as-override
PEx2(config-router-af)#  exit-address-family
PEx2(config-router)# exit
PEx2(config)#no ip vrf CustAB
PEx2(config)#interface serial 0/0.102
PEx2(config-subif)#ip vrf forwarding CustB
PEx2(config-subif)#ip address 150.x.x2.34 255.255.255.240
```

Task 2: Establish CE-PE Connectivity for Central Site Internet Access

Configuration steps on CE*x*1A:

```
CEx1A(config)#interface serial0/0.114 point-to-point
CEx1A(config-subif)#ip add 150.x.x1.66 255.255.255.240
CEx1A(config-subif)#frame-relay interface-dlci 114
CEx1A(config-subif)router ospf 1
CEx1A(config-router)#network 150.x.0.0 0.0.255.255 area 0
CEx1A(config-router)#passive-interface serial0/0.114
```

Configuration steps on CE*x*2B:

```
CEx2B(config)#interface serial0/0.114 point-to-point
CEx2B(config-subif)#ip add 150.x.x2.66 255.255.255.240
CEx2B(config-subif)#frame-relay interface-dlci 114
CEx2B(config-subif)router ospf 2
CEx2B(config-router)#network 150.x.0.0 0.0.255.255 area 0
CEx2B(config-router)#passive-interface serial0/0.114
```

Configuration steps on PE*x*1:

```
PEx1(config)#interface serial0/0.114 point-to-point
PEx1(config-subif)#ip add 150.x.x1.65 255.255.255.240
PEx1(config-subif)#frame-relay interface-dlci 114
PEx1(config-subif)#ip router isis
PEx1(config-subif)#router isis
PEx1(config-router)#passive-interface serial0/0.114
```

Configuration steps on PE*x*2:

```
PEx2(config)#interface serial0/0.114 point-to-point
PEx2(config-subif)#ip add 150.x.x2.65 255.255.255.240
PEx2(config-subif)#frame-relay interface-dlci 114
PEx2(config-subif)#ip router isis
PEx2(config-subif)#router isis
PEx2(config-router)#passive-interface serial0/0.114
```

Task 3: Establish Central Site Connectivity for Internet Access

Configuration steps on CE*x*1A:

```
CEx1A(config)#ip route 10.1.x1.0 255.255.255.0 null0
CEx1A(config)#ip route 10.1.x2.0 255.255.255.0 null0
CEx1A(config)#router bgp 650x1
CEx1A(config-router)#neigh 150.x.x1.65 remote-as 65001
CEx1A(config-router)#network 10.1.x1.0 mask 255.255.255.0
CEx1A(config-router)#network 10.1.x2.0 mask 255.255.255.0
```


Configuration steps on CE*x*2B:

```
CEx2B(config)#ip route 10.2.x1.0 255.255.255.0 null0
CEx2B(config)#ip route 10.2.x2.0 255.255.255.0 null0
CEx2B(config)#router bgp 650x2
CEx2B(config-router)#nei 150.x.x2.65 remote-as 65001
CEx2B(config-router)#network 10.2.x1.0 mask 255.255.255.0
CEx2B(config-router)# network 10.2.x2.0 mask 255.255.255.0
```

Configuration steps on PE*x*1:

```
PEx1(config)#access-list x0 permit 0.0.0.0
PEx1(config)#access-list x0 permit 10.1.x1.0
PEx1(config)#access-list x0 permit 10.1.x2.0
PEx1(config)#route-map Inet_A permit 10
PEx1(config-route-map)#match ip address x0
PEx1(config-route-map)#set extcommunity rt 100:600
PEx1(config-route-map)#exit
PEx1(config)#ip vrf Internet
PEx1(config-vrf)#rd 100:600
PEx1(config-vrf)#route-target import 100:600
PEx1(config-vrf)#export map Inet_A
PEx1(config)#interface serial0/0.114
PEx1(config-subif)#ip vrf forwarding Internet
PEx1(config-subif)#ip address 150.x.x1.65 255.255.255.240
PEx1(config-subif)#exit
PEx1(config)#ip prefix-list DefOnly permit 0.0.0.0/0
PEx1(config)#router bgp 65001
PEx1(config-router)#address-family ipv4 vrf Internet
PEx1(config-router-af)#neighbor 150.x.x1.66 remote 650x1
PEx1(config-router-af)#neighbor 150.x.x1.66 activate
PEx1(config-router-af)#neighbor 150.x.x1.66 prefix-list DefOnly out
```

Configuration steps on PE*x*2:

```
PEx2(config)#access-list x0 permit 0.0.0.0
PEx2(config)#access-list x0 permit 10.2.x1.0
PEx2(config)#access-list x0 permit 10.2.x2.0
PEx2(config)#route-map Inet_B permit 10
PEx2(config-route-map)#match ip address x0
PEx2config-route-map)#set extcommunity rt  100:600
PEx2(config-route-map)#exit
PEx2(config)#ip vrf Internet
PEx2(config-vrf)#rd 100:600
PEx2(config-vrf)#route-target import 100:600
PEx2(config-vrf)#export map Inet_B
PEx2(config)#interface serial0/0.114
PEx2(config-subif)#ip vrf forwarding Internet
PEx2(config-subif)#ip add 150.x.x2.65 255.255.255.240
```

```
PEx2(config-subif)#exit
PEx2(config)#ip prefix-list DefOnly permit 0.0.0.0/0
PEx2(config)#router bgp 65001
PEx2(config-router)#address-family ipv4 vrf Internet
PEx2(config-router-af)#neighbor 150.x.x2.66 remote 650x2
PEx2(config-router-af)#neighbor 150.x.x2.66 activate
PEx2(config-router-af)#neighbor 150.x.x2.66 prefix-list DefOnly out
```

Task 4: Establish Remote Site Internet Connectivity Through the Central Site Router

Configuration steps on the PE routers:

```
PEx1(config)# ip route vrf CustA 0.0.0.0 0.0.0.0 10.1.x1.49
PEx1(config)#router bgp 65001
PEx1(config-router)#address-family ipv4 vrf CustA
PEx1(config-router-af)#redistribute static
PEx1(config-router-af)#default-information originate

PEx2(config)# ip route vrf CustB 0.0.0.0 0.0.0.0 10.2.x2.49
PEx2(config)#router bgp 65001
PEx2(config-router)#address-family ipv4 vrf CustB
PEx2(config-router-af)#redistribute static
PEx2(config-router-af)#default-information originate
```

Lab 8-1 Answer Key: Implementing Basic MPLS TE

When you complete this activity, your router will be similar to the results here, with differences that are specific to your SP.

Task 1: Log the Existing Traffic Flow

Configuration steps on the PE routers:

```
PEx1(config)#mpls ip propagate-ttl
```

```
PEx2(config)#mpls ip propagate-ttl
```

Configuration steps on the P routers:

```
Px1(config)#mpls ip propagate-ttl
```

```
Px1(config)#mpls ip propagate-ttl
```

Task 2: Configure MPLS TE Support on the PE and P Routers

Configuration steps on the PE*x** routers:

```
PEx*(config)#mpls traffic-eng tunnels
PEx*(config)#router isis
PEx*(config-router)#mpls traffic-eng router-id loopback0
PEx*(config-router)#metric-style wide
PEx*(config-router)#mpls traffic-eng level-2

PEx*(config)#interface serial 0/0
PEx*(config)#ip rsvp bandwidth 128
PEx*(config)#interface serial 0/0.111
PEx*(config-subif)#mpls traffic-eng tunnels
PEx*(config-subif)#ip rsvp bandwidth 128
```

Note	When steps show "PE*x**," implement them on both the PE*x*1 and PE*x*2 routers.

Configuration steps on the P*x** routers:

```
Px*(config)#mpls traffic-eng tunnels
Px*(config)#router isis
Px*(config-router)#mpls traffic-eng router-id loopback0
Px*(config-router)#metric-style wide
Px*(config-router)#mpls traffic-eng level-2
```

Note	When steps show "P*x**," implement them on both the P*x*1 and P*x*2 routers.

```
Px1(config)#interface serial 0/0
Px1(config-if)#ip rsvp bandwidth 128
Px1(config-if)#interface serial 0/0.111
Px1(config-subif)#mpls traffic-eng tunnels
Px1(config-subif)#ip rsvp bandwidth 128
Px1(config-subif)#interface serial 0/0.112
Px1(config-subif)#mpls traffic-eng tunnels
Px1(config-subif)#ip rsvp bandwidth 128
Px1(config-subif)#interface serial 0/0.2x1
Px1(config-subif)#ip rsvp bandwidth 128
Px1(config-subif)#mpls traffic-eng tunnels

Px2(config)#interface serial 0/0
Px2(config-if)#ip rsvp bandwidth 128
Px2(config-if)#interface serial 0/0.111
Px2(config-subif)#mpls traffic-eng tunnels
Px2(config-subif)#ip rsvp bandwidth 128
Px2(config-subif)#interface serial 0/0.112
Px2(config-subif)#mpls traffic-eng tunnels
Px2(config-subif)#ip rsvp bandwidth 128
Px2(config-subif)#interface serial 0/0.2x2
Px2(config-subif)#ip rsvp bandwidth 128
Px2(config-subif)#mpls traffic-eng tunnels
```

Task 3: Configure MPLS TE Tunnels

Configuration steps on the PE*x*1 router:

```
PEx1(config)#interface Tunnel0
PEx1(config-if)#ip unnumbered Loopback10
PEx1(config-if)#tunnel destination 192.168.x.33
PEx1(config-if)#tunnel mode mpls traffic-eng
PEx1(config-if)#tunnel mpls traffic-eng autoroute announce
PEx1(config-if)#tunnel mpls traffic-eng priority 0 0
PEx1(config-if)#tunnel mpls traffic-eng bandwidth  100
PEx1(config-if)# tunnel mpls traffic-eng path-option 10 explicit name PEx1_by_P1
PEx1(config-if)#mpls ip
PEx1(config-if)#exit
PEx1(config)#mpls traffic-eng reoptimize timers frequency 30
PEx1(config)#ip explicit-path name PEx1_by_P1
PEx1(cfg-ip-expl-path)#next-address 192.168.x.81
PEx1(cfg-ip-expl-path)#next-address 192.168.100.129
PEx1(cfg-ip-expl-path)#next-address 192.168.x.97
PEx1(cfg-ip-expl-path)#next-address 192.168.x.33
```

Configuration steps on the PE*x*2 router:

```
PEx2(config)#interface Tunnel0
PEx2(config-if)#ip unnumbered Loopback10
PEx2(config-if)#tunnel destination 192.168.x.17
PEx2(config-if)#tunnel mode mpls traffic-eng
PEx2(config-if)#tunnel mpls traffic-eng autoroute announce
PEx2(config-if)#tunnel mpls traffic-eng priority 0 0
PEx2(config-if)#tunnel mpls traffic-eng bandwidth  100
PEx2(config-if)# tunnel mpls traffic-eng path-option 10 explicit name PEx2_by_P1
PEx2(config-if)#mpls ip
PEx2(config-if)#exit
PEx2(config)#mpls traffic-eng reoptimize timers frequency 30
PEx2(config)#ip explicit-path name PEx2_by_P1
PEx2(cfg-ip-expl-path)#next-address 192.168.x.97
PEx2(cfg-ip-expl-path)#next-address 192.168.100.129
PEx2(cfg-ip-expl-path)#next-address 192.168.x.81
PEx2(cfg-ip-expl-path)#next-address 192.168.x.17
```