Implementing Cisco MPLS

Version 2.1

Lab Guide

Text Part Number: ILSG Production Services: 11.18.04

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

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

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: 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: Overlapping VPNs
- Lab 6-2: Merging Service Providers
- Lab 6-3: Common Services VPN
- Lab 7-1: Separate Interface for Internet Connectivity
- Lab 7-2: Multisite Internet Access
- Lab 7-3: Internet Connectivity in an MPLS VPN

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, data-link connection identifier (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, and details of the physical and logical connectivity in the laboratory, and also information about the addressing scheme and IGP routing. The class will be divided into pods (where *x* represents your assigned pod number). Each pod will contain the router types as defined in the table.

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

Router Naming Convention

Router Role	Description
`P (Provider)`	Px1 and Px2 are core routers in the network of the provider.
`PE(Provider Edge)`	PEx1 and PEx2 are edge interfaces routers from provider to customer network.
`CE(Customer Edge)`	CEx1A and CEx2A, and CEx1B and CEx2B are customer edge routers for respective customer A and customer B.

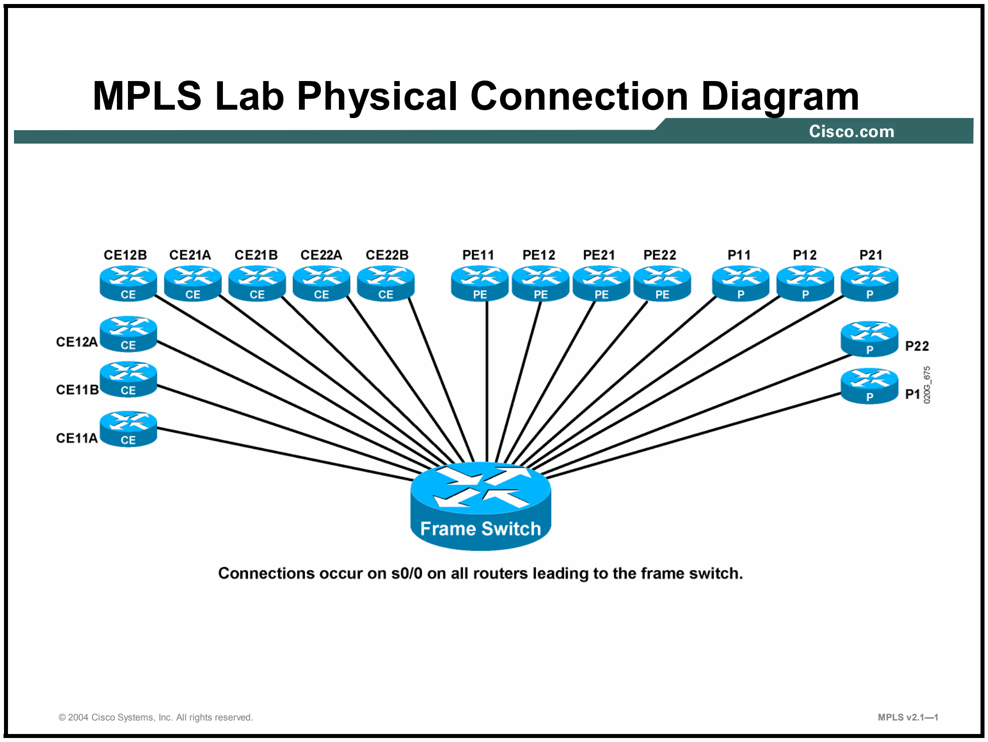

Physical connectivity has been provided by preconfigured permanent virtual circuits (PVCs) defined by their respective DLCIs. The first serial interface of each router (P, PE, and CE) is connected to a Frame Relay switch. The DLCI values for all Frame Relay virtual circuits are shown in the DLCI identification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

DLCI Identification

Source Router Type	Destination Router Type	DLCI
CE*x*1A	PE*x*1	101
CE*x*1B	PE*x*1	102
CE*x*2A	PE*x*2	101
CE*x*2B	PE*x*2	102
PE*x*1	CE*x*1A	101
PE*x*1	CE*x*1B	102
PE*x*1	P*x*1	111
PE*x*2	CE*x*2A	101
PE*x*2	CE*x*2B	102
PE*x*2	P*x*2	111
P*x*1	PE*x*1	111
P*x*1	P*x*2	112
P*x*2	PE*x*2	111
P*x*2	P*x*1	112

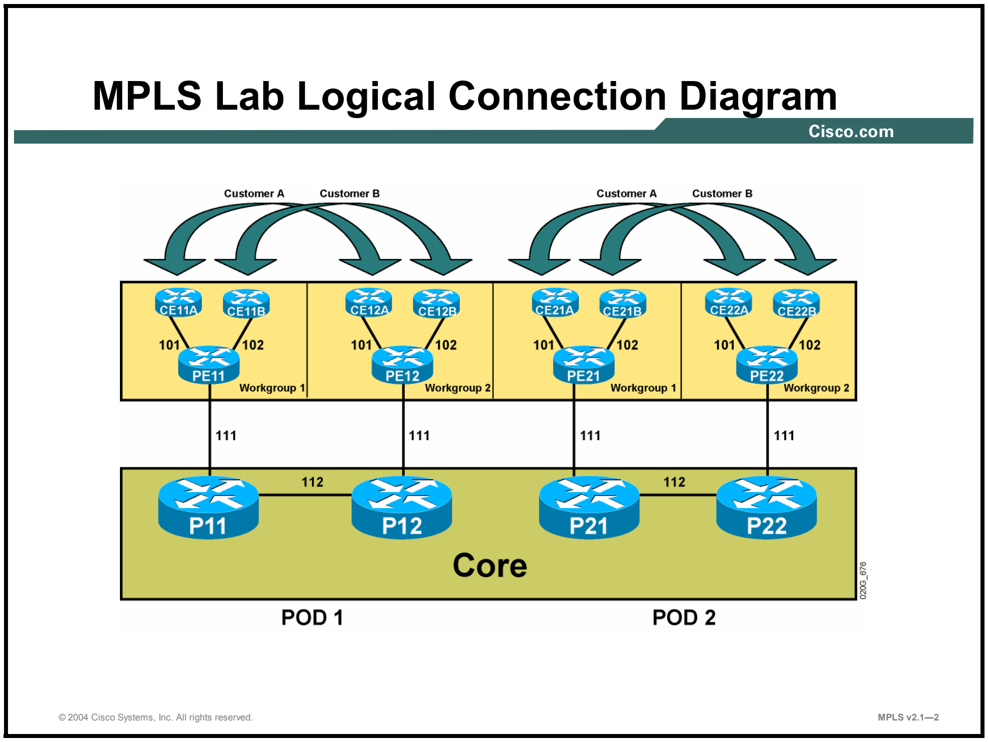

This visual represents the logical connection of each pod. The frame relay DLCI information is included from the DLCI identification table.

Each pod has two P routers creating the core of the service provider network. Each P router connects to the PE router that supports the point of presence (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 pod is further divided into two workgroups. Each workgroup should configure its respective left or right side of the pod. For example, Pod 1 workgroup 1 should configure P11, PE11, CE11A, and CE11B. This leaves workgroup 2 to configure P12, PE12, CE12A, and CE12B.

Your workgroup will still depend on the other workgroup to complete end-to-end connectivity for customer A and customer B. Each customer has a location on each side of the workgroups. An example is customer A with sites CE11A and CE12A. Site CE11A is connected to PE11 with workgroup 11; however, the other site CE12A is connected to the other PE12 router with workgroup 12.

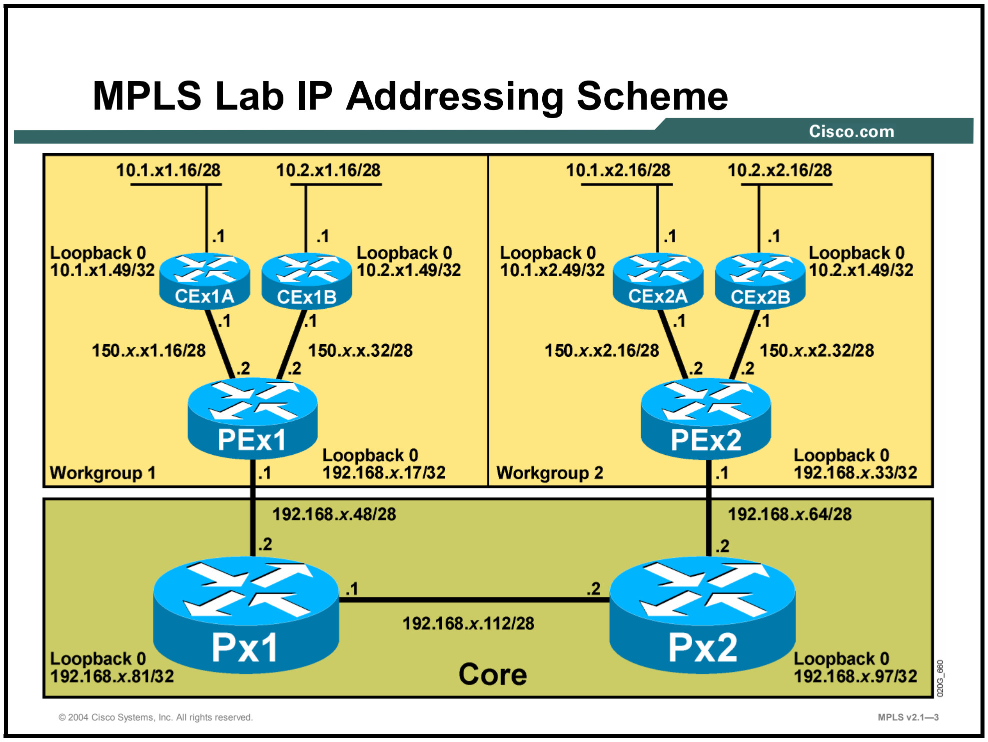

The IP addressing of routers has been performed using the allocations scheme detailed in the IP host address table. Note that *x* equals your pod number.

IP Host Address

Parameter	Value
CEx1A (S0/0.101)	150.x.x1.17/28
CEx1A (loopback0)	10.1.x1.49/32
CEx1A (E0/0)	10.1.x1.17/28
CEx2A (S0/0.101)	150.x.x2.17/28
CEx2A (loopback0)	10.1.x2.49/32
CEx2A (E0/0)	10.1.x2.17/28
CEx1B (S0/0.102)	150.x.x1.33/28
CEx1B (loopback0)	10.2.x1.49/32
CEx1B (E0/0)	10.2.x1.17/28
CEx2B (S0/0.102)	150.x.x2.33/28
CEx2B (loopback0)	10.2.x2.49/32
CEx2B (E0/0)	10.2.x2.17/28
PEx1 (S0/0.101)	150.x.x1.18/28
PEx1 (S0/0.102)	150.x.x1.34/28
PEx1 (loopback0)	192.168.x.17/32
PEx1 (S0/0.111)	192.168.x.49/28
PEx2 (S0/0.101)	150.x.x2.18/28
PEx2 (S0/0.102)	150.x.x2.34/28
PEx2 (loopback0)	192.168.x.33/32
PEx2 (S0/0.111)	192.168.x.65/28
Px1 (S0/0.111)	192.168.x.50/28
Px1 (S0/0.112)	192.168.x.113/28
Px2 (S0/0.111)	192.168.x.66/28
Px2 (S0/0.112)	192.168.x.114/28

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 required to complete this activity:

- Cisco IOS documentation

Command List

The table describes the commands used in this activity.

IP, IGP, and Interface Commands

Command	Description
network *network-number [network-mask]* **no 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* **no 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 Local Management Interface (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. Workgroup 1 and 2 of each pod should configure their respective group of routers.

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

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

Step 3 Configure and enable each customer CE router interface, subinterface, and loopback for appropriate DLCI and IP addressing.

Step 4 Proceed to the activity verification.

Activity Verification

You have completed this task when you attain these results:

- Pinged the remote end of each serial link from each router to verify that each link is operational
- Pinged the loopback interface of a remote router

Task 2: Configuring 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 workgroup 1 and 2 of each pod:

Step 1 On each customer CE router, enable the RIP version 2 (RIPv2) routing process. Disable the auto summary 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 auto summary feature of this routing protocol.

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

Step 4 Proceed to the activity verification.

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 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 150.x.0.0 subnetworks of all P and PE routers are displayed in the IP routing table.

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

When you complete this activity, your router will be similar to the following, with differences that are specific to your pod. The PE routers only need the EIGRP network 150.x.0.0 command for testing. Then remove the network statement. CE routers will need network 150.x.0.0 later in lab 5.1, and you could add the network statement in this lab.

Task 2: Configuring 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: 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
- Disable MPLS TTL propagation
- Configure conditional label distribution

Visual Objective

The figure illustrates what you will accomplish in this activity.

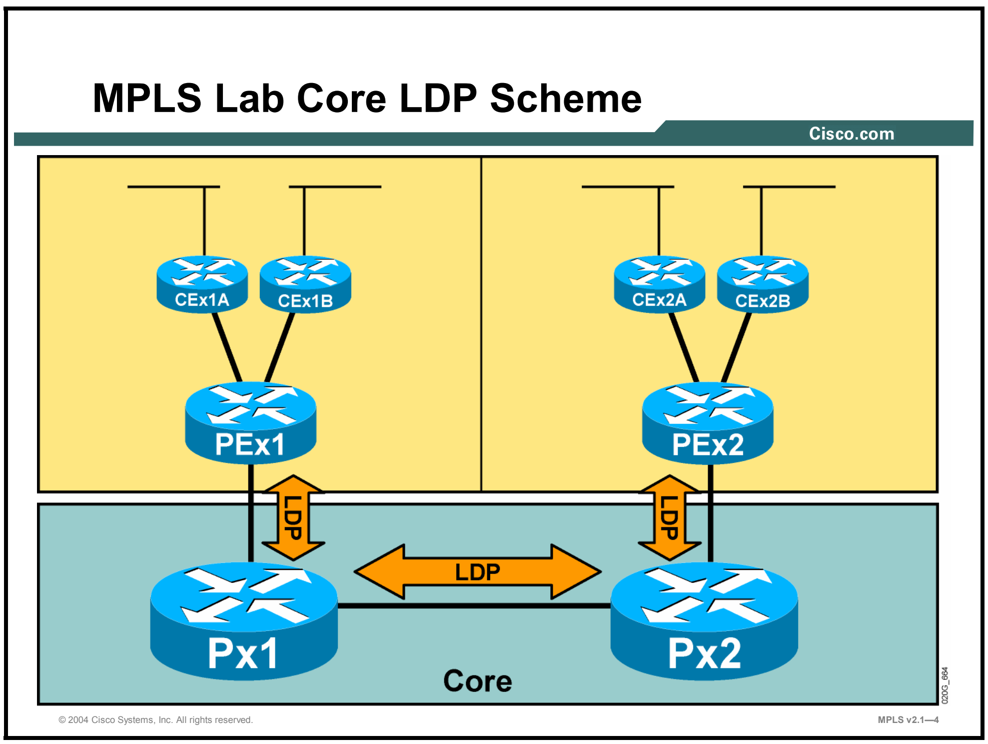

Required Resources

This is the resource required to complete this activity:

- Cisco IOS documentation

Command List

The table describes the commands used in this activity.

MPLS Commands

Command	Description
access-list *access-list-number* {**permit** | **deny**} {*type-code wild-mask* | *address mask*} **no 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 on the RP card, use the **ip cef** command in global configuration mode. To disable CEF, use the **no** form of this command.
mpls ip **no 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 traffic engineering [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 **no 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** } [**no**] **mpls label protocol**	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.

Command	Description
mpls ldp advertise-labels [**for** prefix-access-list [**to** peer-access-list]] **no 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: Enabling 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 workgroup.

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

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

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

```
P11#sh 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 nei
    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, verify that LDP has allocated a label for each prefix in its 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, 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, 8 subnets, 3 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
D       192.168.x.0/24 is a summary, 00:49:20, Null0
C       192.168.x.17/32 is directly connected, Loopback0
     150.x.0.0/16 is variably subnetted, 3 subnets, 2 masks
C       150.x.11.16/28 is directly connected, Serial0/0.101
D       150.x.0.0/16 is a summary, 00:49:20, Null0
C       150.x.11.32/28 is directly connected, Serial0/0.102

Px1#sh mpls ldp bindings
  tib entry: 150.x.0.0/16, rev 16
        local binding:  tag: 20
        remote binding: tsr: 192.168.x.17:0, tag: imp-null
        remote binding: tsr: 192.168.x.97:0, tag: 20
  tib entry: 150.x.11.16/28, rev 18
        remote binding: tsr: 192.168.x.17:0, tag: imp-null
  tib entry: 150.x.11.32/28, rev 19
        remote binding: tsr: 192.168.x.17:0, tag: imp-null
  tib entry: 192.168.x.0/24, rev 17
        remote binding: tsr: 192.168.x.17:0, tag: imp-null
  tib entry: 192.168.x.17/32, rev 14
        local binding:  tag: 19
        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.33/32, rev 10
```

```
        local binding:  tag: 18
        remote binding: tsr: 192.168.x.17:0, tag: 20
        remote binding: tsr: 192.168.x.97:0, tag: 17
  tib entry: 192.168.x.48/28, rev 12
        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: 18
  tib entry: 192.168.x.64/28, rev 6
        local binding:  tag: 17
        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.81/32, rev 8
        local binding:  tag: imp-null
        remote binding: tsr: 192.168.x.17:0, tag: 19
        remote binding: tsr: 192.168.x.97:0, tag: 16
  tib entry: 192.168.x.97/32, rev 2
        local binding:  tag: 16
        remote binding: tsr: 192.168.x.17:0, tag: 16
        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: 17
        remote binding: tsr: 192.168.x.97:0, tag: imp-null
```

- On each of your routers, verify that LDP has received a label of the subnetworks and loopback interfaces of the other core routers.

```
Px1#sh mpls ldp bindings
  tib entry: 150.x.0.0/16, rev 16
        local binding:  tag: 20
        remote binding: tsr: 192.168.x.17:0, tag: imp-null
        remote binding: tsr: 192.168.x.97:0, tag: 20
  tib entry: 150.x.11.16/28, rev 18
        remote binding: tsr: 192.168.x.17:0, tag: imp-null
  tib entry: 150.x.11.32/28, rev 19
        remote binding: tsr: 192.168.x.17:0, tag: imp-null
  tib entry: 192.168.x.0/24, rev 17
        remote binding: tsr: 192.168.x.17:0, tag: imp-null
  tib entry: 192.168.x.17/32, rev 14
        local binding:  tag: 19
        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.33/32, rev 10
        local binding:  tag: 18
        remote binding: tsr: 192.168.x.17:0, tag: 20
        remote binding: tsr: 192.168.x.97:0, tag: 17
  tib entry: 192.168.x.48/28, rev 12
```



```
        local binding:  tag: imp-null
        remote binding: tsr: 192.168.x.17:0, tag: imp-null
        remote binding: tsr: 192.168.1.97:0, tag: 18
  tib entry: 192.168.x.64/28, rev 6
        local binding:  tag: 17
        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.81/32, rev 8
        local binding:  tag: imp-null
        remote binding: tsr: 192.168.x.17:0, tag: 19
        remote binding: tsr: 192.168.x.97:0, tag: 16
  tib entry: 192.168.x.97/32, rev 2
        local binding:  tag: 16
        remote binding: tsr: 192.168.x.17:0, tag: 16
        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: 17
        remote binding: tsr: 192.168.x.97:0, tag: imp-null
```

- Perform a traceroute from your PE router to the loopback address of the PE router of the other workgroup and verify that the results display the associated labels.

```
Tracing the route to 192.168.x.33

  1 192.168.x.50 [MPLS: Label 18 Exp 0] 164 msec 196 msec 200 msec
  2 192.168.x.114 [MPLS: Label 17 Exp 0] 56 msec 56 msec 56 msec
  3 192.168.x.65 40 msec 40 msec
```

Task 2: Disabling TTL Propagation

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

Activity Procedure

Complete these steps:

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

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

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

Activity Verification

You have completed this task when you attain these results:

- You have performed a traceroute from your PE router to the loopback address of the PE router of the other workgroup and compared this display to the display obtained in the previous task.

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

  1 192.168.x.65 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 re-enable TTL propagation. Doing so may affect the results of the traces shown in the lab activity verification because additional hops and labs will be displayed.

Task 3: Configuring Conditional Label Distribution

For the label binding displays that you did in Task 2, 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 wasted label space and resources necessary 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.

Workgroup 1 will configure PE*x*1 and P*x*1. Workgroup 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#sh mpls for
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  point1point
17     Pop tag     192.168.x.112/28  0          Se0/0.111  point1point
18     17          192.168.x.64/28   0          Se0/0.111  point1point
19     Pop tag     192.168.x.81/32   0          Se0/0.111  point1point
20     18          192.168.x.33/32   0          Se0/0.111  point1point
```

Step 2 Note that an LSP has been built to the WAN interface that connects the other PE and P router. 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 workgroup has completed its configuration tasks.

Activity Verification

You have completed this task when you attain these results:

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

```
PE11#sh mpls f
Local  Outgoing    Prefix            Bytes tag  Outgoing   Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
16     16          192.168.1.97/32   0          Se0/0.111  point1point
17     Untagged    192.168.1.112/28  0          Se0/0.111  point1point
18     Untagged    192.168.1.64/28   0          Se0/0.111  point1point
19     Pop tag     192.168.1.81/32   0          Se0/0.111  point1point
20     18          192.168.1.33/32   0          Se0/0.111  point1point
```

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 displayed the LDP bindings.

```
P11#sh mpls ldp bind
  tib entry: 150.x.0.0/16, rev 31
        local binding:  tag: 20
        remote binding: tsr: 192.168.1.97:0, tag: 20
        remote binding: tsr: 192.168.1.17:0, tag: imp-null
  tib entry: 150.x.11.16/28, rev 36
        remote binding: tsr: 192.168.1.17:0, tag: imp-null
  tib entry: 150.x.11.32/28, rev 37
        remote binding: tsr: 192.168.1.17:0, tag: imp-null
  tib entry: 192.168.1.17/32, rev 35
        local binding:  tag: 19
        remote binding: tsr: 192.168.1.97:0, tag: 19
        remote binding: tsr: 192.168.1.17:0, tag: imp-null
  tib entry: 192.168.1.33/32, rev 32
        local binding:  tag: 18
        remote binding: tsr: 192.168.1.97:0, tag: 17
        remote binding: tsr: 192.168.1.17:0, tag: 20
  tib entry: 192.168.1.48/28, rev 26
        local binding:  tag: imp-null
  tib entry: 192.168.1.64/28, rev 27
        local binding:  tag: 17
  tib entry: 192.168.1.81/32, rev 34
        local binding:  tag: imp-null
        remote binding: tsr: 192.168.1.97:0, tag: 16
        remote binding: tsr: 192.168.1.17:0, tag: 19
  tib entry: 192.168.1.97/32, rev 33
        local binding:  tag: 16
        remote binding: tsr: 192.168.1.97:0, tag: imp-null
        remote binding: tsr: 192.168.1.17:0, tag: 16
  tib entry: 192.168.1.112/28, rev 30
        local binding:  tag: imp-null
```

Note The prefix assigned to the WAN interface connecting the other P and PE routers no longer has 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: Removing 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.

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

Activity Procedure

Complete these steps:

Step 1 Remove conditional label distribution.

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

Activity Verification

You have completed this activity when you attain these results:

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

```
PEx1#sh mpls for
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  point1point
17     Pop tag     192.168.x.112/28  0          Se0/0.111  point1point
18     17          192.168.x.64/28   0          Se0/0.111  point1point
19     Pop tag     192.168.x.81/32   0          Se0/0.111  point1point
20     18          192.168.x.33/32   0          Se0/0.111  point1point
```

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

When you complete this activity, your router will be similar to the following, with differences that are specific to your pod.

Task 1: Enabling 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
```

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

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

Task 2: Disabling TTL Propagation

Configuration steps on PE*x*1 and PEx2:

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

Configuration steps on P*x*1 and Px2:

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

Task 3: Configuring Conditional Label Distribution

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

Configuration steps on PE*x*1:

```
PEx1(config)#no tag-switching advertise-tags
PEx1(config)#tag-switching advertise-tags for 90
PEx1(config)#access-list 90 permit 150.x.0.0 0.0.255.255
PEx1(config)#access-list 90 permit 192.168.x.16 0.0.0.15
PEx1(config)#access-list 90 permit 192.168.x.32 0.0.0.15
PEx1(config)#access-list 90 permit 192.168.x.80 0.0.0.15
PEx1(config)#access-list 90 permit 192.168.x.96 0.0.0.15
```

Configuration steps on PE*x*2:

```
PEx2(config)#no tag-switching advertise-tags
PEx2(config)#tag-switching advertise-tags for 90
PEx2(config)#access-list 90 permit 150.x.0.0 0.0.255.255
PEx2(config)#access-list 90 permit 192.168.x.16 0.0.0.15
PEx2(config)#access-list 90 permit 192.168.x.32 0.0.0.15
PEx2(config)#access-list 90 permit 192.168.x.80 0.0.0.15
PEx2(config)#access-list 90 permit 192.168.x.96 0.0.0.15
```

Configuration steps on P*x*1:

```
Px1(config)#no tag-switching advertise-tags
Px1(config)#tag-switching advertise-tags for 90
Px1(config)#access-list 90 permit 150.x.0.0 0.0.255.255
Px1(config)#access-list 90 permit 192.168.x.16 0.0.0.15
Px1(config)#access-list 90 permit 192.168.x.32 0.0.0.15
Px1(config)#access-list 90 permit 192.168.x.80 0.0.0.15
Px1(config)#access-list 90 permit 192.168.x.96 0.0.0.15
```

Configuration steps on P*x*2:

```
Px2(config)#no tag-switching advertise-tags
Px2(config)#tag-switching advertise-tags for 90
Px2(config)#access-list 90 permit 150.x.0.0 0.0.255.255
Px2(config)#access-list 90 permit 192.168.x.16 0.0.0.15
Px2(config)#access-list 90 permit 192.168.x.32 0.0.0.15
Px2(config)#access-list 90 permit 192.168.x.80 0.0.0.15
Px2(config)#access-list 90 permit 192.168.x.96 0.0.0.15
```

Task 4: Removing Conditional Label Distribution

Configuration steps on PE*x*1 and PEx2:

```
PEx*(config)#tag-switching advertise-tags
```

Configuration steps on P*x*1 and Px2:

```
Px*(config)#tag-switching advertise-tags
```


Lab 5-1: 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 pod has been given the task of creating two simple VPNs to support two new customers (customer A and customer B) who 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 workgroup
- 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.

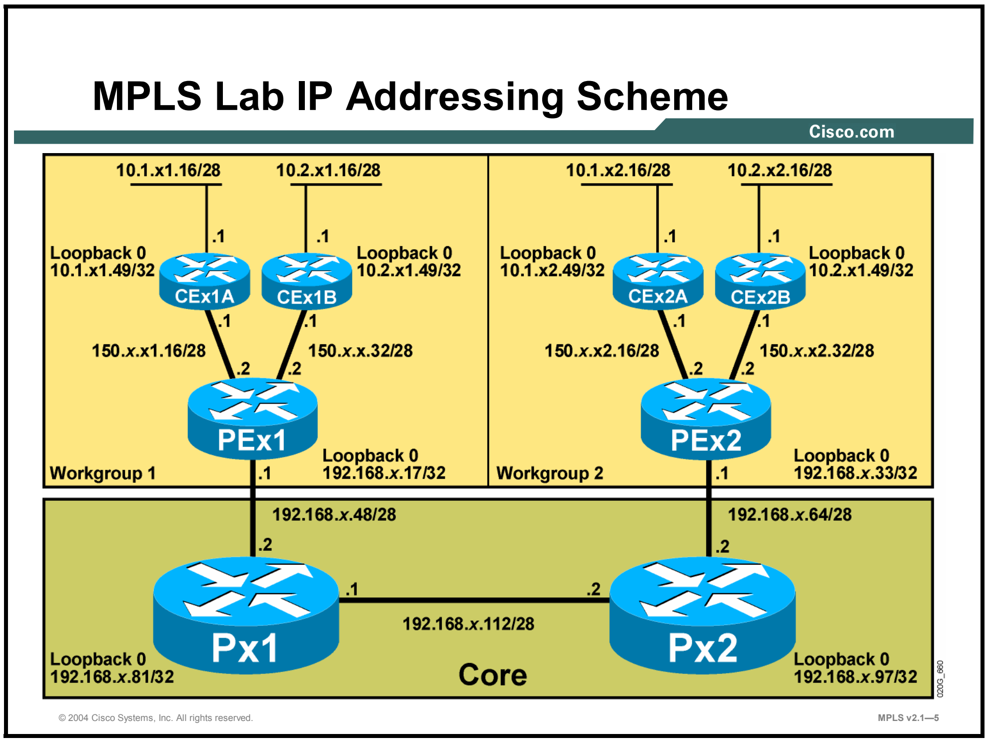

These activities rely on Lab 3-1: Establishing the Core MPLS Environment, in which you established MPLS connectivity in your backbone.

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 workgroup routers or other service providers.

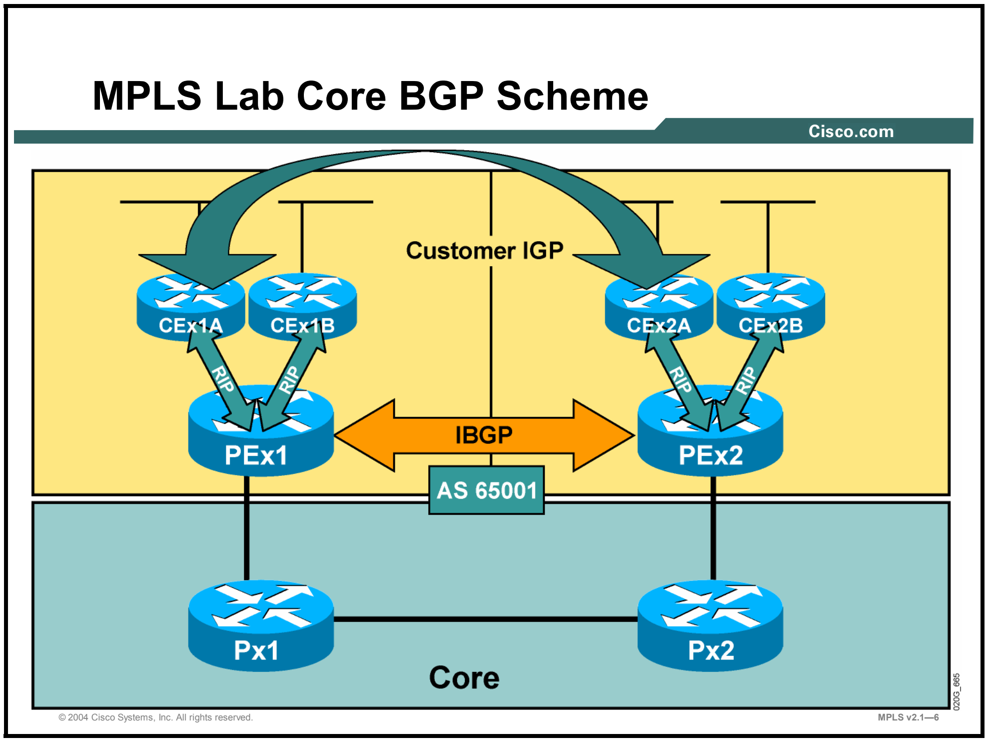

This activity contains tasks that enable you to configure your core MPLS VPN infrastructure and to establish a simple any-to-any VPN service for a customer.

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 required to complete this activity:

- Cisco IOS documentation

Command List

The table describes the commands 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.
neighbor *ip-address* **activate**	Activates an exchange of routes from address family under the configuration for the specified neighbor.

Command	Description
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 communities 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 the 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 a RT to a VRF.
show ip bgp neighbor	Displays information on global BGP neighbors.
show ip bgp vpnv4 vrf *vrf-name*	Displays VPN IPv4 (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.
telnet host /vrf *vrf-name*	Makes a Telnet connection to a CE router connected to the specified VRF.

Task 1: Configuring Multiprotocol BGP

In this section of the activity, you will configure MP-BGP between the PE routers in your workgroup.

Workgroup 1 will configure MP-BGP on PE*x*1, and workgroup 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 auto summary feature.

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

Step 3 Verify that the other workgroup 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.

```
PEx1#sh ip bgp sum

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



PEx2#sh ip bgp sum

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#sh bgp nei

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: Configuring Virtual Routing and Forwarding Tables

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

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, CExA and CExB).

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 with the **address-family ipv4 vrf** v*rf-name* command.

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

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:

- You verified that you have the proper configuration of your VRF tables with the **show ip vrf detail** command. You should get a printout similar to the one here:

```
PEx1#sh ip vrf detail
VRF Customer_A; default RD x:10; default VPNID <not set>
  Interfaces:
    Serial0/0.101
  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 Customer_B; default RD x:20; default VPNID <not set>
  Interfaces:
    Serial0/0.102
  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
```

- Check the routing protocols running in your VRF with the **show ip protocol vrf** command. When executed on PEx1, it will produce a printout similar to the one here:

```
PEx1#sh ip prot vrf Customer_A
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
`    10.0.0.0
    150.x.0.0
  Routing Information Sources:
    Gateway         Distance      Last Update
    150.x.x1.17          120      00:00:27
  Distance: (default is 120)

PEx1#sh ip prot vrf Customer_B
```

```
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
    10.0.0.0
    150.x.0.0
  Routing Information Sources:
    Gateway         Distance      Last Update
    150.x.x1.33          120      00:00:07
  Distance: (default is 120)
```

- Verify the per-VRF routing table on the PE router with the **show ip route vrf** command. It will produce a printout similar to the one here:

```
PEx1#sh ip route vrf Customer_A
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#sh ip route vrf Customer_B
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
```

- Use the **show ip bgp vpnv4 vrf** command to display the BGP routing table associated with a VRF. The printout from the PEx1 router is shown here:

```
PEx1#show ip bgp vpnv4 vrf Customer_A
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 Customer_A)
*> 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 Customer_B
```

```
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 Customer_B)
*> 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 ?
```

- On a CE router, use the **show ip route** command to verify that the router is receiving all VPN routes. Also verify that no routes from the other customer or the MPLS core are being received. On CE*x*1A, the printout is similar to the one here:

```
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, 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

Use 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 = 144/146/148 ms
```

- Use the **show ip route** command on the PE routers to verify that the customer routes are not in the global 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, 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
```

- Use **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:
.....
```

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

```
PEx1#ping vrf Customer_A 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 Customer_B 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 5-1 Answer Key: Initial MPLS VPN Setup

When you complete this activity, your router will be similar to the following, with differences that are specific to your pod.

Task 1: Configuring Multiprotocol 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: Configuring Virtual Routing and Forwarding Tables

Configuration steps on PE*x*1:

```
PEx1(config)#ip vrf Customer_A
PEx1(config-vrf)#rd x:10
PEx1(config-vrf)#route-target both x:10
PEx1(config)#ip vrf Customer_B
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 Customer_A
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 Customer_B
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 Customer_A
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 Customer_B
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 Customer_A
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 Customer_B
PEx1(config-router-af)#no auto-summary
PEx1(config-router-af)#redistribute rip
```

Configuration steps on PEx2:

```
PEx2(config)#ip vrf Customer_A
PEx2(config-vrf)#rd x:10
PEx2(config-vrf)#route-target both x:10
PEx2(config)#ip vrf Customer_B
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 Customer_A
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 Customer_B
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 Customer_A
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 Customer_B
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 Customer_A
PEx2(config-router-af)#no auto-summary
PEx2(config-router-af)#redistribute rip
PEx2(config-router)#address-family ipv4 vrf Customer_B
PEx2(config-router-af)#no auto-summary
PEx2(config-router-af)#redistribute rip
```


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 remain running RIP as the IGP.

Visual Objective

The figure illustrates what you will accomplish in this activity.

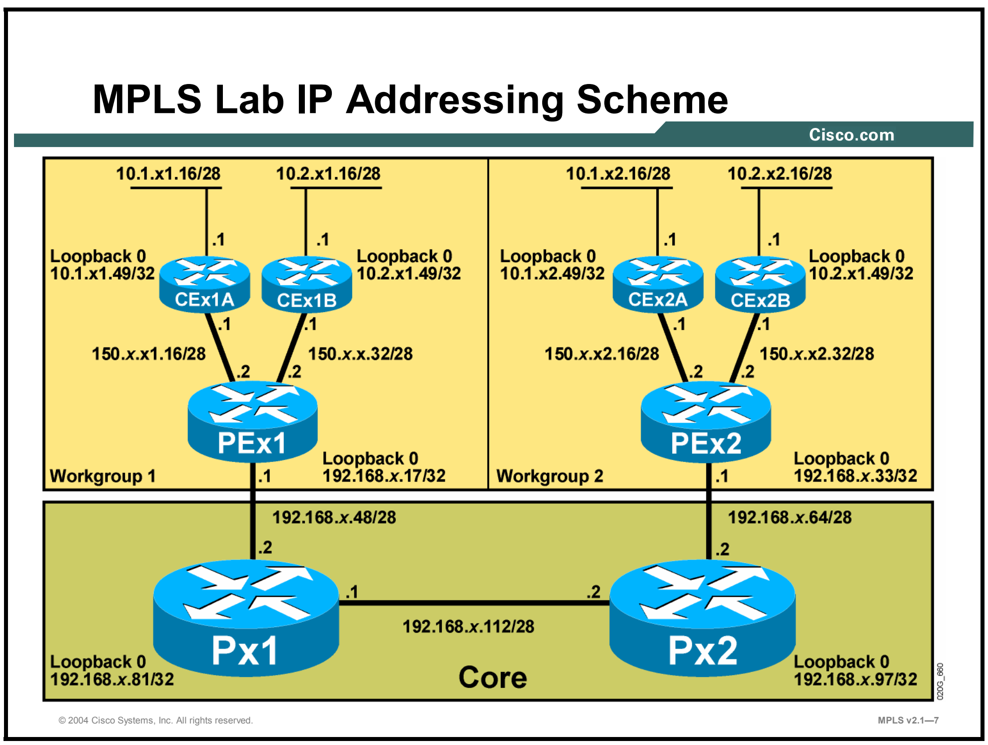

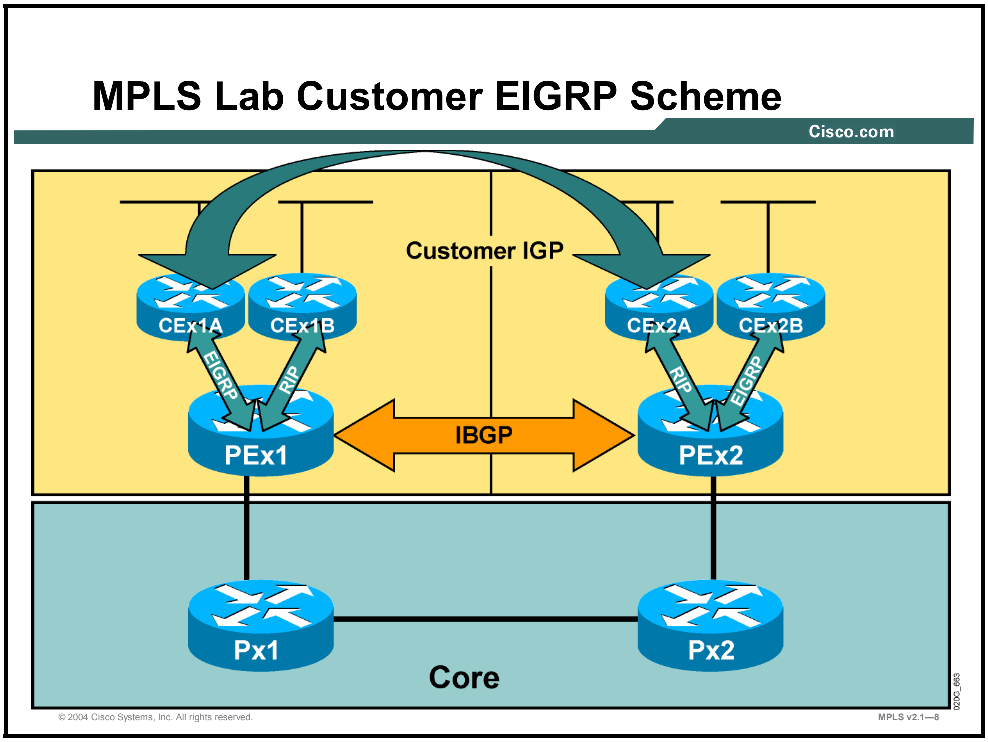

Required Resources

This is the resource required to complete this activity:

- Cisco IOS documentation

Command List

The table describes the commands used in this activity.

OSPF Commands

Command	Description
address-family ipv4 [**multicast** | **unicast** | **vrf** *vrf-name*]	Enters address family configuration mode and creates a VRF. The VRF name (or tag) must match the VRF name that was created in Step 3 from Task 2.
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: Enabling an EIGRP VPN

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

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

Each workgroup 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. Workgroup 1 will configure CE*x*1A, and workgroup 2 will configure CE*x*2B. Use your *x#* 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.

Note Do not forget to remove the address family from the RIP routing process. This action will disable the sites still running RIP as the CE-PE routing protocol.

Step 2 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 3 On your assigned PE router, configure redistribution of BGP into EIRGP with the **address-family ipv4 vrf** *vrf-name* command Disable the auto summary 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#sh ip eigrp int

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#sh ip eigrp vrf Customer_A nei
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#sh ip eigrp vrf Customer_B nei
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
```

- Check the EIGRP topology database on the CE routers.

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

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

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.x.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)
```

- Verify 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 150.x.x2.33
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.x.x2.33, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 144/147/152 ms

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 = 144/147/152 ms

CEx1B#trace 150.x.x2.33
Type escape sequence to abort.
Tracing the route to 150.x.x2.33
1 150.x.x1.34 12 msec 12 msec 12 msec
2 150.x.x2.34 64 msec 60 msec 60 msec
3 150.x.x2.33 77 msec 76 msec *

CEx1A#trace 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 64 msec 60 msec 64 msec
3 150.x.x2.17 76 msec 76 msec *

PEx1#ping vrf Customer_A 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

PEx2#ping vrf Customer_A 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
```

```
PEx1#trace vrf Customer_B 10.2.x2.49
Type escape sequence to abort.
Tracing the route to 10.2.x2.49
  1 150.x.x2.33 60 msec 60 msec *

PEx2#trace vrf Customer_A 10.1.x1.49
Type escape sequence to abort.
Tracing the route to 10.1.x1.49
  1 150.x.x1.17 60 msec 60 msec *
```

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

When you complete this activity, your router will be similar to the following, with differences that are specific to your pod.

Task 1: Enabling 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 Customer_A
PEx1(config)#router eigrp 1
PEx1(config-router)#address-family ipv4 vrf Customer_A
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 Customer_A
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 Customer_B
PEx2(config-router)#router eigrp 1
PEx2(config-router)#address-family ipv4 vrf Customer_B
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 Customer_B
PEx2(config-router-af)#no redistribute rip
PEx2(config-router-af)#redistribute eigrp x metric 1
```


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 complete 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
- Complete the OSPF migration

Visual Objective

The figure illustrates what you will accomplish in this activity.

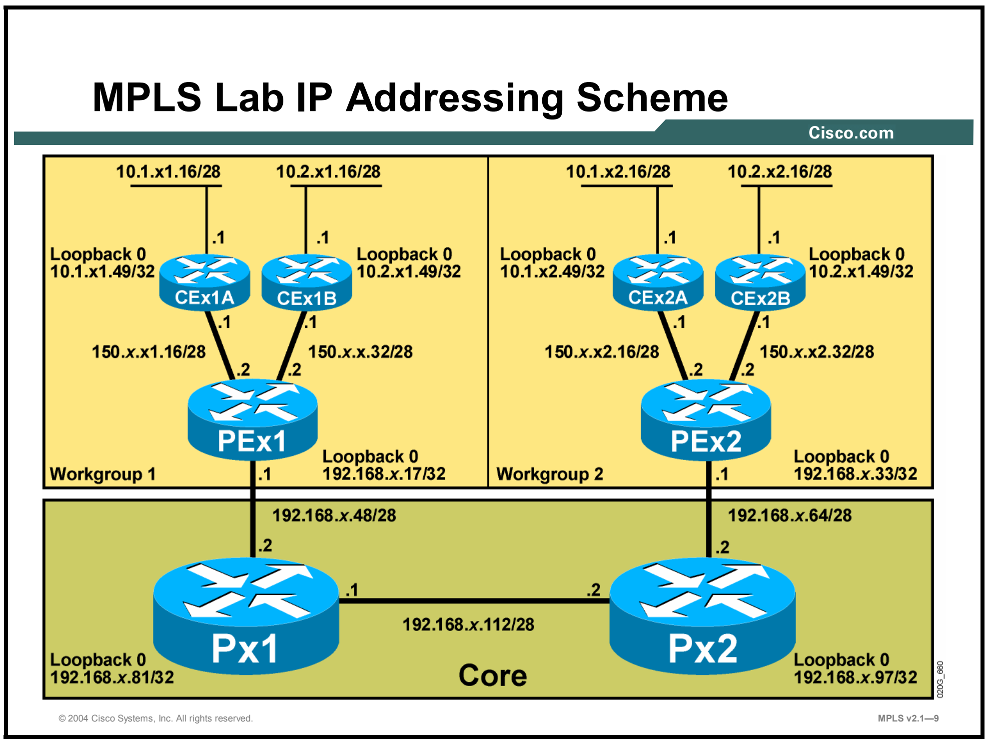

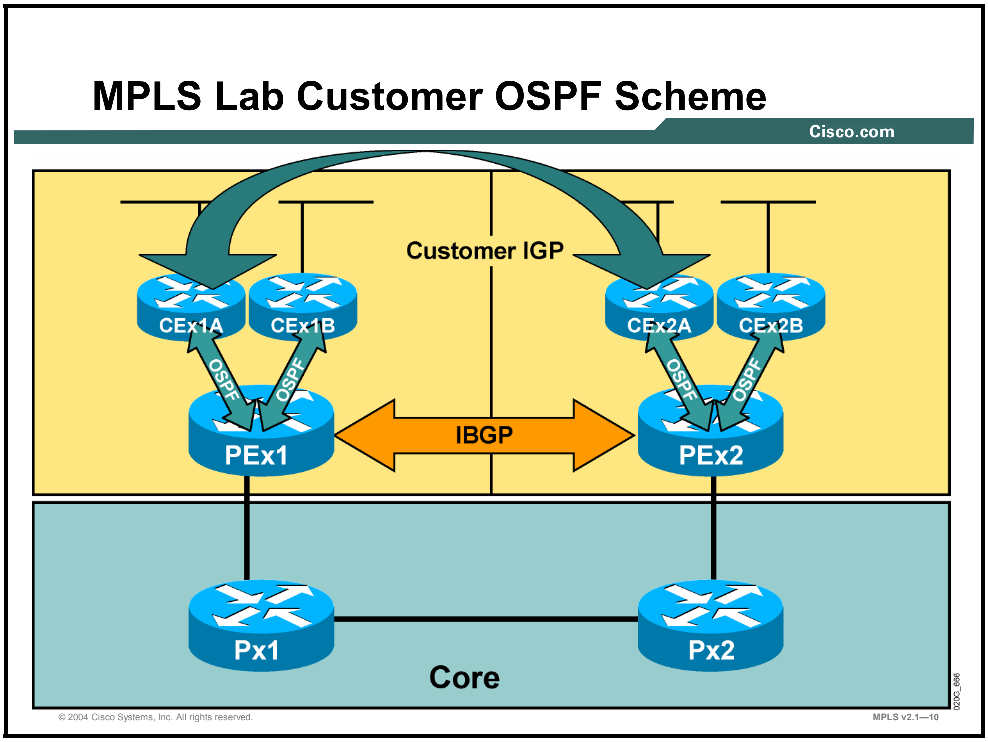

Required Resources

This is the resource required to complete this activity:

- Cisco IOS documentation

Command List

The table describes the commands 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: Configuring OSPF as the PE-CE Routing Protocol

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

Each workgroup 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 (use an OSPF process ID of 1 for workgroup 1 and a process ID of 2 for workgroup 2) areas 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 workgroup 1 and a process ID of 2 for workgroup 2) in the VRFs on 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#sh ip ospf nei
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#sh ip ospf nei

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
```

- Check 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*x1*A is shown here:

```
CEx1A#sh ip ospf data

           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        0x80000005 0x007C30 3
150.x.x1.18     150.x.x1.18     216         0x80000004 0x000E87 2

                Summary Net Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum
10.1.x1.16      10.1.x1.49      1744        0x80000002 0x0012C1
10.1.x2.16      150.x.x1.18     1186        0x80000001 0x00CDD7
10.1.x2.49      150.x.x1.18     1186        0x80000001 0x0082FB
150.x.x2.16     150.x.x1.18     1186        0x80000001 0x00CD94

                Router Link States (Area 1)

Link ID         ADV Router      Age         Seq#       Checksum Link count
10.1.x1.49      10.1.x1.49      1744        0x80000002 0x00532E 1

                Summary Net Link States (Area 1)

Link ID         ADV Router      Age         Seq#       Checksum
10.1.x1.49      10.1.x1.49      1744        0x80000002 0x00C6E5
10.1.x2.16      10.1.x1.49      1294        0x80000001 0x000E45
10.1.x2.49      10.1.x1.49      1294        0x80000001 0x00C269
150.x.x1.16     10.1.x1.49      1853        0x80000002 0x000D04
150.x.x2.16     10.1.x1.49      1294        0x80000001 0x000E02

                Summary ASB Link States (Area 1)

Link ID         ADV Router      Age         Seq#       Checksum
150.x.x1.18     10.1.x1.49      332         0x80000002 0x0045B9
```

- Check the IP routing table on CE*x*1A and note the OSPF interarea (IA) routes 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, 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
```

- Verify 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 Customer_B 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 Customer_A 10.1.x2.49
Type escape sequence to abort.
Tracing the route to 10.1.x2.49

  1 150.x.x1.17 80 msec 100 msec *

PEx1#trace vrf Customer_B 10.2.x1.49
Type escape sequence to abort.
Tracing the route to 10.2.x1.49

  1 150.x.x1.33 60 msec 60 msec *
```

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

When you complete this activity, your router will be similar to the following, with differences that are specific to your pod.

Task 1: Configuring 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 Customer_B
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 Customer_B
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 Customer_A
PEx1(config)#router ospf 1 vrf Customer_A
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 Customer_A
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 Customer_A
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 Customer_A
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 Customer_B
PEx2(config)#router ospf 2 vrf Customer_B
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 Customer_B
PEx2(config-router-af)#no redistribute eigrp x
PEx2(config-router-af)#redistribute ospf 2
```

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 do this conversion in a phased implementation. 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.

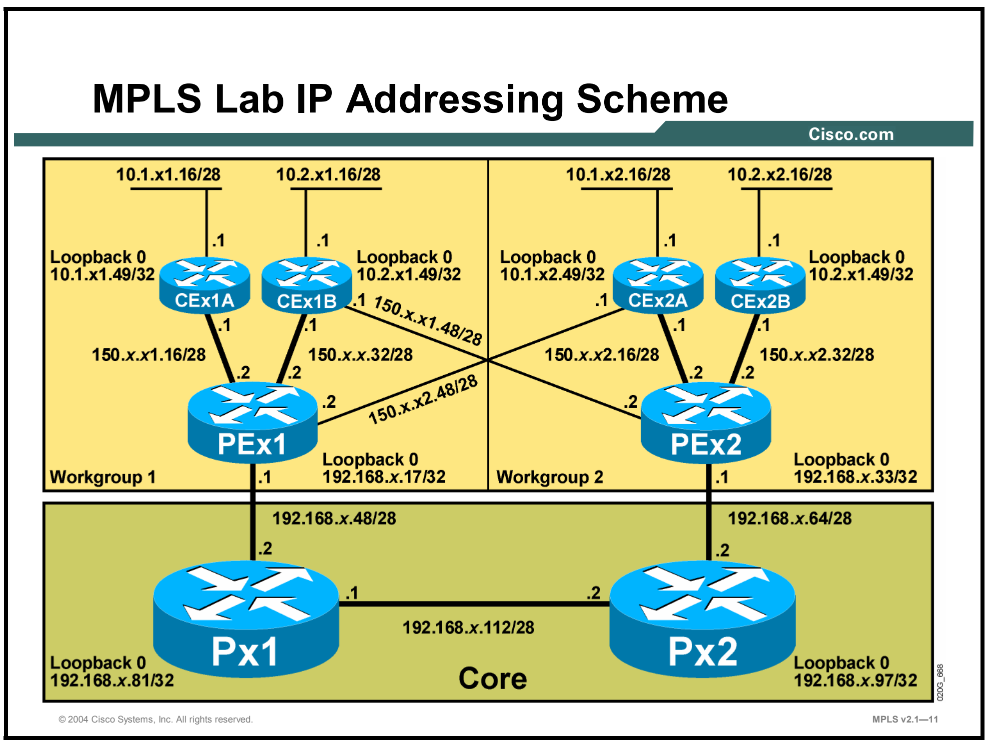

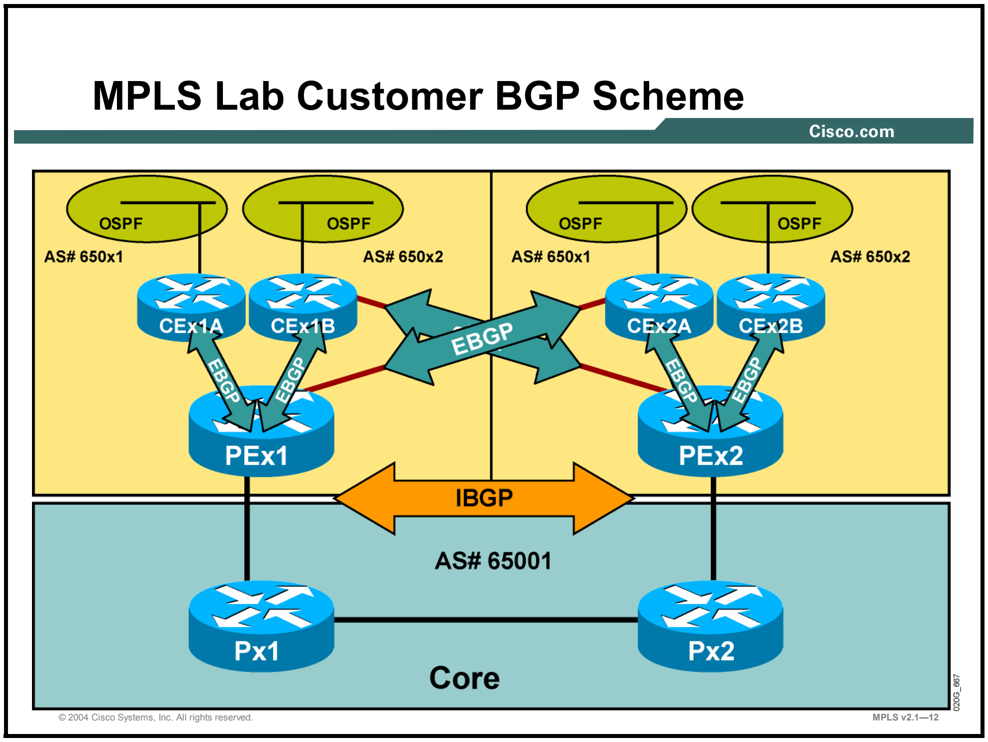

Required Resources

This is the resource required to complete this activity:

- Cisco IOS documentation

Command List

The table describes the commands 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: Configuring 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. Workgroup 1 will convert customer A (CE*x*1A and CE*x*2A), and workgroup 2 will convert customer B (CE*x*1B and CE*x*2B) to establish a simple VPN. Each workgroup 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 AS650x1 for customer A and AS 650x2 for customer B. Disable the auto summary 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 of your customer sites are using the same AS number, 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#sh ip bgp sum
BGP router identifier 10.1.x1.49, local AS number 650x1
BGP table version is 10, main routing table version 10
9 network entries and 9 paths using 1197 bytes of memory
2 BGP path attribute entries using 120 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 activity 9/30 prefixes, 9/0 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#sh ip bgp
BGP table version is 63, 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#sh 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: 1:10 (default for vrf Customer_A)
*> 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: 1:20 (default for vrf Customer_B)
*> 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: Configuring the Backup PE-CE Link

In this task, you will enable the backup links on the PE routers. Workgroup 1 will establish the link between its PEx1 router and the CEx2A router, and workgroup 2 will establish the link between its PEx2 router and the CEx1B 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.

Step 2 Add the backup link to the appropriate VRF.

Which VRF is CE*x*1B added to?

____

Which VRF is CE*x*2A added to?

____

Step 3 Configure IP addresses and data-link connection identifiers (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.x1.49/28*	113	PE*x*1	150.x.x1.50/28	113
CE*x*1B	150.*x.x2.49/28*	113	PE*x*2	150x.x2.50/28	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 Customer_B 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 = 28/28/32 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 Customer_A 150.x.x1.49
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.x.x1.49, timeout is 2 seconds:
!!!!!
```

- Check BGP connectivity with the **show ip bgp summary** command on the CE routers.

```
CEx2A#sh ip bgp sum
BGP router identifier 10.1.x2.49, local AS number 650x2
BGP table version is 10, main routing table version 10
9 network entries and 9 paths using 1197 bytes of memory
2 BGP path attribute entries using 120 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 activity 9/30 prefixes, 9/0 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#sh 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#sh ip bgp vpn all
BGP table version is 36, local router ID is 192.168.1.17
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 Customer_A)
*> 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: 1:20 (default for vrf Customer_B)
```



```
* 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 ?
```

Task 3: Selecting 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 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.

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#sh 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    150.x.x2.50                    50      0 65001 65001 ?
```

```
*>                  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 ?
```

- Verify 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#sh 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: 1:10 (default for vrf Customer_A)
*>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 ?
*>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: 1:20 (default for vrf Customer_B)
*>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 ?
```

- Shut down the link from the local PE router to the CE router.

- Verify that the backup link (the link to your local PE router) is being used. Use the **show ip bgp** command to verify this.

```
CEx1B#sh ip bgp
BGP table version is 14, 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    150.x.x2.50                    50      0 65001 65001 ?
*>                  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.x1.32/28   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   150.x.x2.50                    50      0 65001 65001 ?
*>                  0.0.0.0                  0         32768 ?
```

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

```
CEx1B#sh 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    150.x.x2.50                    50      0 65001 65001 ?
*>                  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 5-4 Answer Key: Running BGP Between PE and CE Routers

When you complete this activity, your router will be similar to the following, with differences that are specific to your pod.

Task 1: Configuring 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:

```
!******* Workgroup 1 *************
PEx1(config)#router bgp 65001
PEx1(config-router)#address-family ipv4 vrf Customer_A
PEx1(config-router-af)#no redistribute ospf 1
PEx1(config)#no router ospf 1 vrf Customer_A
PEx1(config)#router bgp 65001
```



```
PEx1(config-router)#address-family ipv4 vrf Customer_A
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
!******* Workgroup 2 **************
PEx1(config)#router bgp 65001
PEx1(config-router-af)#address-family ipv4 vrf Customer_B
PEx1(config-router-af)#no redistribute ospf 2
PEx1(config)#no router ospf 2 vrf Customer_B
PEx1(config)#router bgp 65001
PEx1(config-router-af)#address-family ipv4 vrf Customer_B
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:

```
!******* Workgroup 1 *******
PEx2(config)#router bgp 65001
PEx2(config-router)#address-family ipv4 vrf Customer_A
PEx2(config-router-af)#no redistribute ospf 1
PEx2(config)#no router ospf 1 vrf Customer_A
PEx2(config)#router bgp 65001
PEx2(config-router)#address-family ipv4 vrf Customer_A
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
!******* Workgroup 2 *******
PEx2(config-router-af)#address-family ipv4 vrf Customer_B
PEx2(config-router-af)#no redistribute ospf 2
PEx2(config)#no router ospf 2 vrf Customer_B
PEx2(config)#router bgp 65001
PEx2(config-router-af)#address-family ipv4 vrf Customer_B
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: Configuring the Backup PE-CE Link

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 Customer_B
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 Customer_B
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 Customer_A
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 Customer_A
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: Selecting 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: 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 have met 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.

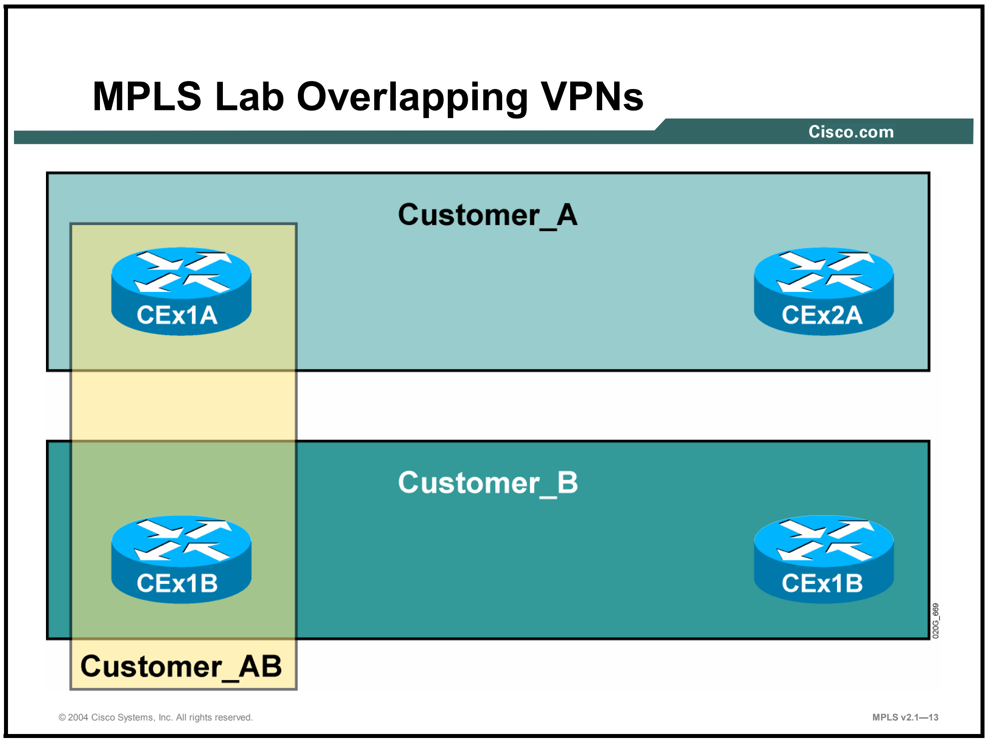

In this lab activity, you will establish overlapping VPNs with the following connectivity goals:

- 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 (Customer_AB):
  - CE*x*1A and CE*x*2B can communicate.

Required Resources

This is the resource required to complete this activity:

- Cisco IOS documentation

Command List

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

Task 1: Designing 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 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 Customer_AB VPN. Coordinate the value of this RT with the other workgroup within your pod.

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 Customer_AB VPN.

Activity Verification

You have completed this task when you attain this result:

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

Task 2: Removing 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 router.

Step 2 Check any other references to CE*x*1A and CE*x*2B in their PE router configuration 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#sh ip vrf int
Interface            IP-Address      VRF                           Protocol
Interface            IP-Address      VRF                           Protocol
Serial0/0.113        150.x.x1.50     Customer_A                    up
Serial0/0.102        150.x.x1.34     Customer_B                    up

PEx2#sh ip vrf int
Interface            IP-Address      VRF                           Protocol
Serial0/0.101        150.x.x2.18     Customer_A                    up
Serial0/0.113        150.x.x2.50     Customer_B                    up
```

- Verify 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#sh ip bgp vpnv4 vrf Customer_A sum
BGP router identifier 192.168.x.17, local AS number 65001
BGP table version is 34, main routing table version 34
7 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#sh ip bgp vpnv4 vrf Customer_B sum
BGP router identifier 192.168.x.33, local AS number 65001
BGP table version is 33, main routing table version 33
5 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: Configuring 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#sh ip vrf int
Interface              IP-Address      VRF                Protocol
Serial0/0.113          150.x.x1.50     Customer_A         up
Serial0/0.101          150.x.x1.18     Customer_AB        up
Serial0/0.102          150.x.x1.34     Customer_B         up

PEx2#sh ip vrf int
Interface              IP-Address      VRF                Protocol
Serial0/0.101          150.x.x2.18     Customer_A         up
Serial0/0.102          150.x.x2.34     Customer_AB        up
Serial0/0.113          150.x.x2.50     Customer_B         up
```

- Verify the BGP neighbors on the PE router with the **show ip bgp vpnv4 vrf summary** command. This 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#sh ip bgp vpnv4 vrf Customer_AB sum
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
7 BGP path attribute entries using 1510 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 2438 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#sh ip bgp vpnv4 vrf Customer_AB sum
BGP router identifier 192.168.x.33, local AS number 65001
BGP table version is 56, main routing table version 56
8 network entries using 968 bytes of memory
8 path entries using 512 bytes of memory
7 BGP path attribute entries using 1510 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 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
```

- Check 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 650x1 in the AS path, and routes announced by CE*x*2B should have 650x2 in the AS path.

```
PEx1#sh ip bgp vpnv4 vrf Customer_AB
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:1001 (default for vrf Customer_AB)
*> 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#sh ip bgp vpnv4 vrf Customer_AB
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: 1:21 (default for vrf Customer_AB)
*>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 ?
```

- Connect to CE*x*1A and perform 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 = 52/58/68 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.x2.33 [AS 650x2] 72 msec 77 msec *
```

- 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 6-1 Answer Key: Overlapping VPNs

When you complete this activity, your router will be similar to the following, with differences that are specific to your pod.

Task 1: Designing Your VPN Solution

Note No configuration steps are required for this task.

Task 2: Removing 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 Customer_A
PEx1(config-router-af)#no neighbor 150.x.x1.17
PEx1(config-vrf)#interface serial0/0.101
PEx1(config-subif)#no ip vrf forwarding Customer_A
```

Note After removing the interface from the VRF, the following message will appear: "% Interface Serial0/0.101 IP address 150.x.*x1*.18 removed due to disabling VRF Customer_A."

Configuration steps on PE*x*2:

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

Note After removing the interface from the VRF, the following message will appear: "% Interface Serial0/0.102 IP address 150.x.*x2*.34 removed due to disabling VRF Customer_B."

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

Note	RDs and RTs listed in these results may or may not match what you have used in this lab task.

Configuration steps on PE*x*1:

```
PEx1(config)#ip vrf Central_AB
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 Central_AB
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 Central_AB
PEx1(config-router-af)#neighbor 150.x.x1.17 remote-as 650x1
PEx1(config-router-af)#neighbor 150.x.x1.17 activate
```

Configuration steps on PE*x2*:

```
PEx2(config)#ip vrf Central_AB
PEx2(config-vrf)#rd x:21
PEx2(config-vrf)#route-target both x:20
PEx2(config-vrf)#route-target both x:1001
PEx1(config-vrf)#interface serial0/0.102
PEx2(config-subif)#ip vrf forwarding Central_AB
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 Central_AB
PEx2(config-router-af)#neighbor 150.x.x2.33 remote-as 650x2
PEx2(config-router-af)#neighbor 150.x.x2.33 activate
```


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 Px1 and Px2 router to P1. In addition, the core Interior Gateway Protocol (IGP) is being converted from Enhanced Interior Gateway Routing Protocol (EIGRP) to Intermediate System-to-Intermediate System (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:

- 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

Workgroup 1 will configure PE*x*1 and Px1, and workgroup 2 will configure PE*x*2 and P*x*2. P1 has been preconfigured.

The figure illustrates what you will accomplish in this activity.

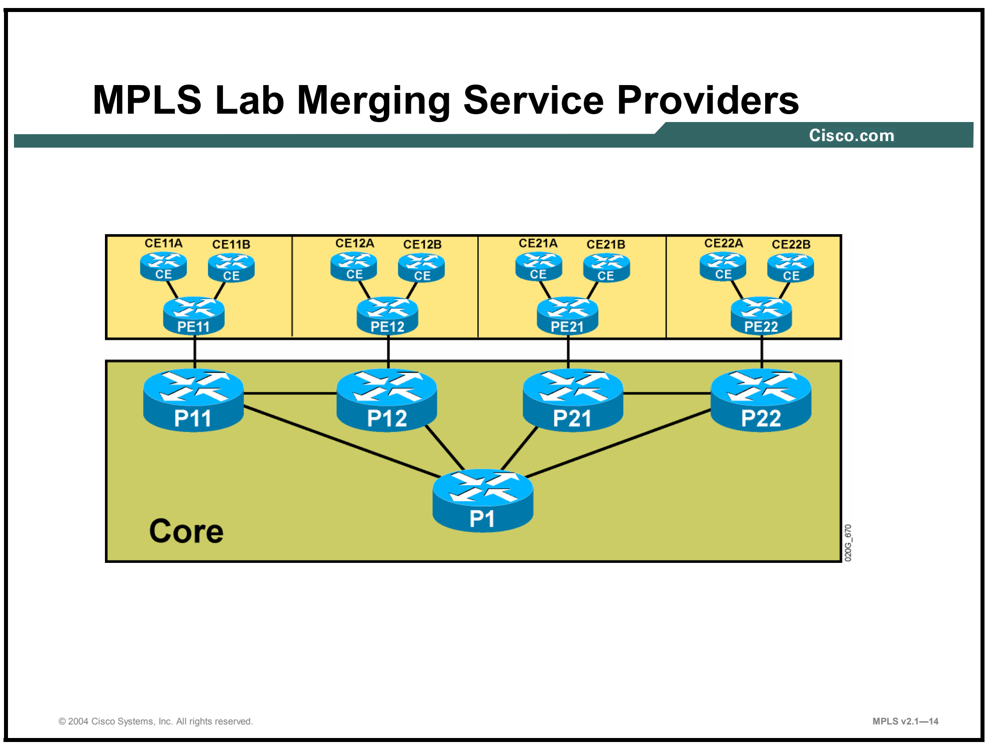

Required Resources

This is the resource required to complete this activity:

- Cisco IOS documentation

Command List

The table describes the commands 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, length, and value objects (TLVs), use the **metric-style wide** command in router configuration mode. To disable this function, use the **no** form of this command.

Task 1: Enabling 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 Label Distribution Protocol (LDP) connectivity between the two routers.

Activity Procedure

Complete these steps:

Step 1 Configure IP addresses and data-link connection identifiers (DLCIs) on this interface using the parameters in the table here.

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

IP Address and DLCI Configuration Parameters

Router	Subinterface	DLCI	IP Address
P11	S0/0.211	211	192.168.100.10/29
P12	S0/0.212	212	192.168.100.18/29
P21	S0/0.221	221	192.168.100.26/29
P22	S0/0.222	222	192.168.100.34/29

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: Migrating 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 workgroup is responsible for the migration of all of your assigned routers. Workgroup 1 will migrate PE*x*1 and P*x*1. Workgroup 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 49.0001.0000.0000.01x1.00	Where x = the POD number
PE*x*2	net 49.0001.0000.0000.01x2.00	
P*x*1	net 49.0001.0000.0000.02x1.00	
P*x*2	net 49.0001.0000.0000.02x2.00	

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 protocol** command to verify that IS-IS is active and enabled on all appropriate interfaces.

```
PEx1#sh ip prot
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
    192.168.1.33
  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
    192.168.100.10       115      00:09:05
    192.168.x.97         115      00:07:27
    192.168.x.113        115      00:09:27
    192.168.x.114        115      00:07:42
    192.168.x.81         115      00:07:37
    192.168.x.33         115      00:07:37
    192.168.x.50         115      00:09:32
    192.168.100.129      115      00:09:05
  Distance: (default is 115)
Px1#sh ip prot
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      00:02:20
  192.168.x.114        115      00:14:40
  192.168.100.18       115      00:14:35
  192.168.x.33         115      00:14:35
  192.168.x.17         115      00:02:20
  192.168.100.129      115      00:02:20
Distance: (default is 115)
```

- Use the **show ip route** command and verify that all routers are sending and receiving the appropriate prefixes.

```
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, 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.50, Serial0/0.111
C       192.168.1.48/28 is directly connected, Serial0/0.111
C       192.168.x.17/32 is directly connected, Loopback0
     192.168.100.0/24 is variably subnetted, 3 subnets, 2 masks
i L2    192.168.100.8/29 [115/20] via 192.168.x.50, Serial0/0.111
i L2    192.168.100.16/29 [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

  Px1#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, 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.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
     192.168.100.0/24 is variably subnetted, 3 subnets, 2 masks
C       192.168.100.8/29 is directly connected, Serial0/0.211
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    192.168.100.129/32 [115/20] via 192.168.100.9, Serial0/0.2x1
```

Task 3: Enabling 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.

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#sh mpls ldp nei
    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.81.11047
        State: Oper; Msgs sent/rcvd: 48/46; Downstream
        Up time: 00:27:52
        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.97.11021 - 192.168.x.81.646
            State: Oper; Msgs sent/rcvd: 47/48; Downstream
            Up time: 00:26:14
            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   192.168.100.18
```

- On your PE router, verify that labels are being received from the other workgroups.

```
PEx1#sh mpls forw
Local  Outgoing    Prefix            Bytes tag  Outgoing   Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
16     Pop tag     192.168.x.81/32   0          Se0/0.111  point2point
17     Pop tag     192.168.x.112/28  0          Se0/0.111  point2point
18     Pop tag     192.168.100.8/29  0          Se0/0.111  point2point
19     17          192.168.x.97/32   0          Se0/0.111  point2point
20     18          192.168.x.64/28   0          Se0/0.111  point2point
22     20          192.168.x.33/32   0          Se0/0.111  point2point
23     Untagged    10.1.x1.16/28[V]  0          Se0/0.101  point2point
24     Untagged    10.1.x1.49/32[V]  0          Se0/0.101  point2point
25     Aggregate   150.x.x1.16/28[V] 2212
30     21          192.168.100.16/29 0          Se0/0.111  point2point
31     22          192.168.100.129/32   \
                                     0          Se0/0.111  point2point
35     Untagged    10.2.x1.16/28[V]  0          Se0/0.102  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
```

Task 4: Enabling IBGP Connectivity for All PE Routers

At this point, you have established LDP connectivity for all of the P routers in your new service provider environment, but you have not yet established BGP connectivity. You now need to establish Internal Border Gateway Protocol (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 clients 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.

Workgroup 1 will configure PE*x*1, and workgroup 2 will configure PE*x*2.

Activity Procedure

Complete these steps:

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

Step 2 Activate your PE router as a client of P1.

Activity Verification

You have completed this task when you attain these results:

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

```
PEx1#sh ip bgp sum
BGP router identifier 192.168.x.17, local AS number 65001
BGP table version is 4, main routing table version 4
1 network entries using 101 bytes of memory
1 path entries using 48 bytes of memory
9 BGP path attribute entries using 540 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 857 total bytes of memory
BGP activity 61/33 prefixes, 98/66 paths, scan interval 60 secs

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

PE11#sh ip bgp nei
BGP neighbor is 150.x.x1.17,  vrf A_Central,  remote AS 65011, external link
  BGP version 4, remote router ID 10.1.x1.49
  BGP state = Established, up for 04:26:36
  Last read 00:00:36, 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
  Message statistics:
    InQ depth is 0
    OutQ depth is 0
                         Sent       Rcvd
    Opens:                  1          1
    Notifications:          0          0
    Updates:               11          1
    Keepalives:           269        269
```



```
    Route Refresh:          0          0
    Total:                281        271
  Default minimum time between advertisement runs is 30 seconds

 For address family: VPNv4 Unicast
  Translates address family IPv4 Unicast for VRF A_Central
  BGP table version 198, neighbor version 198
  Index 1, Offset 0, Mask 0x2
                               Sent       Rcvd
  Prefix activity:             ----       ----
    Prefixes Current:             7          3 (Consumes 384 bytes)
    Prefixes Total:              29          3
    Implicit Withdraw:            4          0
    Explicit Withdraw:           18          0
    Used as bestpath:           n/a          6
    Used as multipath:          n/a          0

                                 Outbound    Inbound
  Local Policy Denied Prefixes:    --------    -------
    Bestpath from this peer:              3        n/a
    Total:                                3          0
  Number of NLRIs in the update sent: max 6, min 0

  Connections established 1; dropped 0
  Last reset never
Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Local host: 150.x.x1.18, Local port: 11005
Foreign host: 150.x.x1.17, Foreign port: 179

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

Event Timers (current time is 0x12EF17C):
Timer          Starts    Wakeups            Next
Retrans           281          0             0x0
TimeWait            0          0             0x0
AckHold           270        215             0x0
SendWnd             0          0             0x0
KeepAlive           0          0             0x0
GiveUp              0          0             0x0
PmtuAger            0          0             0x0
DeadWait            0          0             0x0

iss: 3658226700  snduna: 3658232505  sndnxt: 3658232505     sndwnd:  15040
irs: 3372805470  rcvnxt: 3372810690  rcvwnd:      15567  delrcvwnd:    817

SRTT: 303 ms, RTTO: 330 ms, RTV: 27 ms, KRTT: 0 ms
```

```
minRTT: 40 ms, maxRTT: 352 ms, ACK hold: 200 ms
Flags: higher precedence, nagle

Datagrams (max data segment is 1460 bytes):
Rcvd: 432 (out of order: 0), with data: 270, total data bytes: 5219
Sent: 500 (retransmit: 0, fastretransmit: 0), with data: 280, total data
bytes:4

BGP neighbor is 150.x.x1.33,  vrf Customer_B,  remote AS 65012, external link
  BGP version 4, remote router ID 10.2.x1.49
  BGP state = Established, up for 05:25:54
  Last read 00:00:53, 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
  Message statistics:
    InQ depth is 0
    OutQ depth is 0
                         Sent       Rcvd
    Opens:                  2          2
    Notifications:          0          0
    Updates:                9         16
    Keepalives:           334        334
    Route Refresh:          0          0
    Total:                345        352
  Default minimum time between advertisement runs is 30 seconds

 For address family: VPNv4 Unicast
  Translates address family IPv4 Unicast for VRF Customer_B
  BGP table version 198, neighbor version 198
  Index 2, Offset 0, Mask 0x4
  Overrides the neighbor AS with my AS before sending updates
                                 Sent       Rcvd
  Prefix activity:               ----       ----
    Prefixes Current:               3          4 (Consumes 256 bytes)
    Prefixes Total:                12          4
    Implicit Withdraw:              0          0
    Explicit Withdraw:              9          0
    Used as bestpath:             n/a          4
    Used as multipath:            n/a          0

                                   Outbound    Inbound
  Local Policy Denied Prefixes:    --------    -------
    AS_PATH loop:                       n/a          9
    Bestpath from this peer:              7        n/a
    Total:                                7          9
  Number of NLRIs in the update sent: max 3, min 0
```



```
  Connections established 2; dropped 1
  Last reset 05:26:35, due to Peer closed the session
Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Local host: 150.x.x1.34, Local port: 11003
Foreign host: 150.x.x1.33, Foreign port: 179

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

Event Timers (current time is 0x12F0188):
Timer          Starts    Wakeups            Next
Retrans           336          0             0x0
TimeWait            0          0             0x0
AckHold           337        262             0x0
SendWnd             0          0             0x0
KeepAlive           0          0             0x0
GiveUp              0          0             0x0
PmtuAger            0          0             0x0
DeadWait            0          0             0x0

iss:  488391326  snduna:  488397950  sndnxt:  488397950     sndwnd:  15623
irs: 2402946008  rcvnxt: 2402952696  rcvwnd:      15583  delrcvwnd:    801

SRTT: 301 ms, RTTO: 314 ms, RTV: 13 ms, KRTT: 0 ms
minRTT: 32 ms, maxRTT: 340 ms, ACK hold: 200 ms
Flags: higher precedence, nagle

Datagrams (max data segment is 1460 bytes):
Rcvd: 539 (out of order: 0), with data: 337, total data bytes: 6687
Sent: 605 (retransmit: 0, fastretransmit: 0), with data: 335, total data
bytes:3

BGP neighbor is 150.x.x1.49,  vrf Customer_A,  remote AS 65011, external link
  BGP version 4, remote router ID 10.1.12.49
  BGP state = Established, up for 05:24:42
  Last read 00:00:43, 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
  Message statistics:
    InQ depth is 0
    OutQ depth is 0
                         Sent       Rcvd
    Opens:                  2          2
    Notifications:          0          0
    Updates:               15         12
    Keepalives:           334        334
```

```
    Route Refresh:          0          0
    Total:                351        348
  Default minimum time between advertisement runs is 30 seconds

 For address family: VPNv4 Unicast
  Translates address family IPv4 Unicast for VRF Customer_A
  BGP table version 198, neighbor version 198
  Index 3, Offset 0, Mask 0x8
  Overrides the neighbor AS with my AS before sending updates
                               Sent       Rcvd
  Prefix activity:             ----       ----
    Prefixes Current:             7          4 (Consumes 256 bytes)
    Prefixes Total:              22          4
    Implicit Withdraw:            4          0
    Explicit Withdraw:           11          0
    Used as bestpath:           n/a          0
    Used as multipath:          n/a          0

                                 Outbound    Inbound
  Local Policy Denied Prefixes:    --------    -------
    AS_PATH loop:                      n/a         12
    Bestpath from this peer:            16        n/a
    Total:                              16         12
  Number of NLRIs in the update sent: max 4, min 0

  Connections established 2; dropped 1
  Last reset 05:25:35, due to Peer closed the session
Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Local host: 150.x.x1.50, Local port: 179
Foreign host: 150.x.x1.49, Foreign port: 11001

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

Event Timers (current time is 0x12F115C):
Timer          Starts    Wakeups            Next
Retrans           338          0             0x0
TimeWait            0          0             0x0
AckHold           335        160             0x0
SendWnd             0          0             0x0
KeepAlive           0          0             0x0
GiveUp              0          0             0x0
PmtuAger            0          0             0x0
DeadWait            0          0             0x0

iss: 1459130543  snduna: 1459137294  sndnxt: 1459137294     sndwnd:  15499
irs: 2671381495  rcvnxt: 2671388204  rcvwnd:      15538  delrcvwnd:    846
```

```
SRTT: 305 ms, RTTO: 340 ms, RTV: 35 ms, KRTT: 0 ms
minRTT: 40 ms, maxRTT: 372 ms, ACK hold: 200 ms
Flags: passive open, nagle, gen tcbs

Datagrams (max data segment is 1460 bytes):
Rcvd: 561 (out of order: 0), with data: 335, total data bytes: 6708
Sent: 503 (retransmit: 0, fastretransmit: 0), with data: 337, total data
bytes:0

BGP neighbor is 192.168.100.129,  remote AS 65001, internal link
  BGP version 4, remote router ID 207.69.43.1
  BGP state = Established, up for 00:05:21
  Last read 00:00:21, hold time is 180, keepalive interval is 60 seconds
  Neighbor capabilities:
    Route refresh: advertised and received(old & new)
    Address family IPv4 Unicast: advertised and received
    Address family VPNv4 Unicast: advertised and received
  Message statistics:
    InQ depth is 0
    OutQ depth is 0
                         Sent       Rcvd
    Opens:                  2          2
    Notifications:          0          0
    Updates:                4          6
    Keepalives:            11         11
    Route Refresh:          0          0
    Total:                 17         19
  Default minimum time between advertisement runs is 5 seconds

 For address family: IPv4 Unicast
  BGP table version 4, neighbor version 4
  Index 1, Offset 0, Mask 0x2
                                 Sent       Rcvd
  Prefix activity:               ----       ----
    Prefixes Current:               0          1 (Consumes 48 bytes)
    Prefixes Total:                 0          1
    Implicit Withdraw:              0          0
    Explicit Withdraw:              0          0
    Used as bestpath:             n/a          1
    Used as multipath:            n/a          0

                                   Outbound    Inbound
  Local Policy Denied Prefixes:    --------    -------
    Bestpath from this peer:              1        n/a
    Total:                                1          0
  Number of NLRIs in the update sent: max 0, min 0
```

```
 For address family: VPNv4 Unicast
  BGP table version 198, neighbor version 198
  Index 4, Offset 0, Mask 0x10
  NEXT_HOP is always this router
  Community attribute sent to this neighbor
                               Sent       Rcvd
  Prefix activity:             ----       ----
    Prefixes Current:             7          7 (Consumes 1088 bytes)
    Prefixes Total:              11          7
    Implicit Withdraw:            0          0
    Explicit Withdraw:            4          0
    Used as bestpath:           n/a         17
    Used as multipath:          n/a          0

                                 Outbound    Inbound
  Local Policy Denied Prefixes:    --------    -------
    VPN Imported prefix:                 10        n/a
    Bestpath from this peer:              7        n/a
    Total:                               17          0
  Number of NLRIs in the update sent: max 4, min 0

  Connections established 2; dropped 1
  Last reset 00:05:51, due to Address family activated
Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Local host: 192.168.x.17, Local port: 11006
Foreign host: 192.168.100.129, Foreign port: 179

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

Event Timers (current time is 0x12F2334):
Timer          Starts    Wakeups            Next
Retrans            11          0             0x0
TimeWait            0          0             0x0
AckHold            11          9             0x0
SendWnd             0          0             0x0
KeepAlive           0          0             0x0
GiveUp              0          0             0x0
PmtuAger            0          0             0x0
DeadWait            0          0             0x0

iss: 1034013014  snduna: 1034013734  sndnxt: 1034013734     sndwnd:  16215
irs:  841715741  rcvnxt:  841716609  rcvwnd:      16151  delrcvwnd:    233

SRTT: 232 ms, RTTO: 775 ms, RTV: 543 ms, KRTT: 0 ms
minRTT: 60 ms, maxRTT: 308 ms, ACK hold: 200 ms
```

```
Flags: higher precedence, nagle

Datagrams (max data segment is 536 bytes):
Rcvd: 21 (out of order: 0), with data: 11, total data bytes: 867
Sent: 22 (retransmit: 0, fastretransmit: 0), with data: 10, total data bytes:
79
```

- Verify 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#sh ip bgp vpnv4 vrf Customer_A
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 Customer_A)
*> 10.1.x1.16/28    150.x.x1.17              0             0 65011 ?
*> 10.1.x1.49/32    150.x.x1.17              0             0 65011 ?
*>i10.1.x2.16/28    192.168.x.33             0    100      0 65011 ?
*                   150.x.x1.49            200             0 65011 ?
*>i10.1.x2.49/32    192.168.x.33             0    100      0 65011 ?
*                   150.x.x1.49            200             0 65011 ?
*> 150.x.x1.16/28   150.x.x1.17              0             0 65011 ?
r>i150.x.x1.48/28   192.168.x.33             0    100      0 65011 ?
r                   150.x.x1.49            200             0 65011 ?
*>i150.x.x2.16/28   192.168.x.33             0    100      0 65011 ?
*                   150.x.x1.49            200             0 65011 ?
```

- Verify 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#sh ip route vrf Customer_A

Routing Table: Customer_A
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 (Central_AB), 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 (Central_AB), 04:51:01
B       10.1.x2.49/32 [200/0] via 192.168.x.33, 00:27:42
     150.1.0.0/28 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 (Central_AB), 04:51:01
C       150.x.x1.48 is directly connected, Serial0/0.113

PEx1#sh ip route vrf Customer_B

Routing Table: Customer_B
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
```

Lab 6-2 Answer Key: Merging Service Providers

When you complete this activity, your router will be similar to the following, with differences that are specific to your pod.

Task 1: Enabling Connectivity with the Central P Router

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

Note	These step are for the P router of the pod 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: Migrating the Core to IS-IS

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 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 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
```

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 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 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: Enabling MPLS LDP Connectivity with the Central P Router

Note The subinterface number and DLCI number in the following configurations will match with 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: Enabling IBGP Connectivity for All PE Routers

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: 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 service provider 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:

- Design a network management VPN
- Establish connectivity between the management VRF and customer VRFs by configuring proper route targets

Visual Objective

The figure illustrates what you will accomplish in this activity.

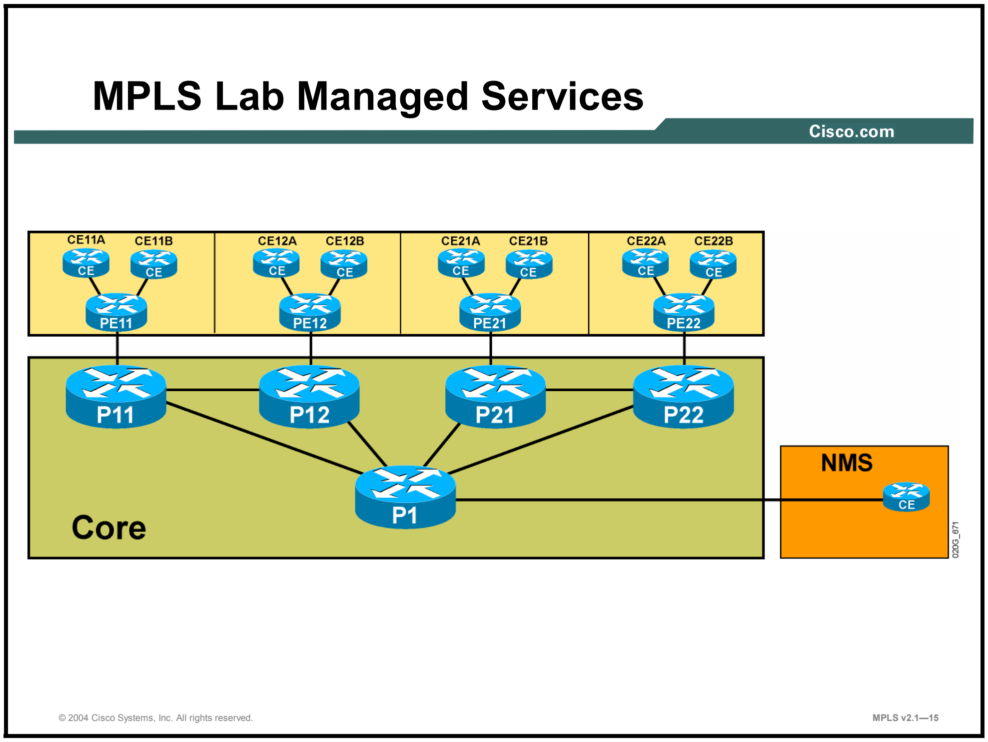

Note	The NMS routers are shared between workgroups and are not configurable.

Required Resources

This is the resource required to complete this activity:

- Cisco IOS documentation

Command List

The table describes the commands used in this activity.

Network Management VPN Commands

Command	Description
export map *name*	Specifies a VRF export route map.
ip prefix-list *name* **permit** *address mask* **ge** *len*	Creates an IP prefix list that matches all prefixes in a specified address space with a subnet mask longer or equal to the specified value.
match ip address prefix-list *list*	Matches a prefix in a route map with a specified IP prefix list.
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: Establishing 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	The following configuration resides on the P1 router and, in this exercise, serves as a 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	If you were implementing a common services VPN from scratch, you would need to configure the supporting PE router using the VRF and routing commands used in previous exercises. In this implementation, the NMS VPN is already configured on the central service PE router, so you will need only to configure the VRF of your customer 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 routes toward the CE router loopback addresses when the addresses are exported from the customer VRF. You also need to import routes toward 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 prefix 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 route-map** command.

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

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 the one here:

```
PEx1#sh ip bgp vpnv4 vrf Customer_A 10.1.x1.49
BGP routing table entry for 1:10:10.1.x1.49/32, version 46
Paths: (1 available, best #1, table Customer_A)
  Advertised to non peer-group peers:
  150.x.x1.49
  650x1, imported path from 1:11:10.1.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:1:10 RT:1:1001 RT:101:501
```

- Using an **extended ping** command, verify 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).
- Using an **extended ping** command, verify 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).
- Verify that your CE router is seeing only prefixes within your VPN and that no prefixes are being leaked from other VPNs.

```
PEx1#sh ip bgp vpnv4 vrf Customer_A
BGP table version is 53, 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 Customer_A)
*> 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 ?
*>i10.10.10.49/32   192.168.100.129          0    100      0 ?
*> 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 ?
```

Lab 6-3 Answer Key: Common Services VPN

When you complete this activity, your router will be similar to the following, with differences that are specific to your pod.

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

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

```
PEx1(config)#ip vrf Customer_A
PEx1(config-vrf)#export map NMS_Cus_A
PEx1(config-vrf)#route-target import 101:500
PEx1(config)#ip vrf A_Central
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 access-list 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 Customer_A
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 Customer_B
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 Customer_B
PEx2(config-vrf)#export map NMS_Cus_B
PEx2(config-vrf)#route-target import 101:500
PEx2(config)#ip vrf B_Central
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: Separate Interface for Internet Connectivity

In many cases, customers may want to retain the traditional Internet access model with a firewall between the customer VPN and the global Internet. This request is usually implemented by using dedicated VPN and Internet subinterfaces on the physical PE-CE link.

Activity Objective

In this activity, you will implement a separate interface for Internet access. After completing this activity, you will be able to meet these objectives:

- Establish CE-PE connectivity for Internet access
- Establish routing between the customer and the Internet

Visual Objective

You will configure additional virtual links (emphasized in the visual) between the central site CE routers (CE*x*1A and CE*x*2B) and their PE routers. These circuits will be in the global routing table, and you will configure static routing between the PE and CE routers. The remote sites (CE*x*1B and CE*x*2A) will access the Internet using the MPLS VPN connection back to its respective central site and then through the newly created link.

The figure illustrates what you will accomplish in this activity.

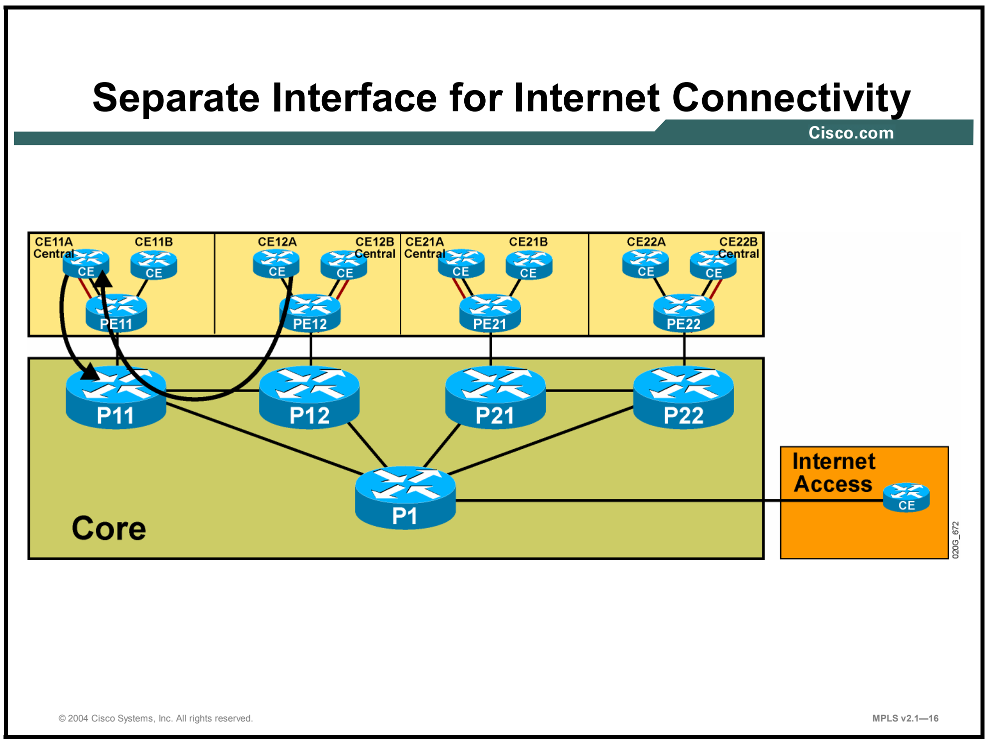

Note	In this lab, the customer addressing scheme is in the private addressing range. In an actual implementation, a NAT service would need to be provided at the customer interface to the Internet access point. Because NAT is outside the scope of this course, this function is omitted, and the lab has been set up to ensure that the customer addressing does not overlap.

Required Resources

This is the resource required to complete this activity:

- Cisco IOS documentation

Command List

The table describes the commands used in this activity.

Internet Access Commands

Command	Description
ip route *prefix mask* **null 0**	Creates a summary route in the IP routing table.

Task 1: Establishing CE-PE Connectivity for Internet Access

In this task, you will add a new subinterface to support Internet access 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 from below.

Router ID	IP Address	DLCI
CEx1A	150.x.x1.66/28	114
CEx2B	150.x.x2.66/28	114

Step 2 Activate the new interface in the Interior Gateway Protocol (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
PEx1	150.x.x1.65/28	114
PEx2	150.x.x2.65/28	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#sh ip int 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 ************************


PEx1#sh ip int 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 ************************
```


Task 2: Establishing Routing Between the Customer and the Internet

In this solution, the customer and the service provider have decided to use static routing for the PE-CE Internet routing protocol. In this task, you will enable a static default route on the CE router that points to the Internet and a static route on the PE router that points to the customer address range.

Activity Procedure

Complete these steps:

Step 1 On the PE router that is supporting your CE router, create a static route that points to the customer address range.

Note Your first choice for the static route would most likely be 10.1.0.0/16 for customer A and 10.2.0.0/16 for customer B. However, if you examine the addressing scheme used in these labs, you will notice that customer A on all pods uses the same 10.1.0.0 address range. The same is true for customer B, which uses 10.2.0.0 on all pods. To ensure that your static routes do not overlap with the other pods, you will need a statement for each customer site.

Step 2 Redistribute this route into BGP so that it will be advertised to the Internet access point.

Step 3 On your CE router, create a default route that will point all unknown routes to the Internet interface.

Step 4 This static route will be used by both the local central sites and the remote VPN sites. Because of this shared use, you will need to interject the route into both the local and remote routing tables. You can accomplish this task by adding a network statement to the BGP process that enables network 0.0.0.0.

Note For security reasons, the customer never wants packets that originate in its network or that are addressed to its network to be sent out to the Internet. Creating a default route that points all unroutable customer packets to the null interface will address this issue.

Activity Verification

You have completed this task when you attain these results:

- You have verified the static route on the PE router.

```
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, 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

**** output omitted *******

     10.0.0.0/24 is subnetted, 2 subnets
S       10.1.x1.0 [1/0] via 150.x.11.66
S       10.1.x2.0 [1/0] via 150.x.11.66

***** output omitted ******



Verify the static routes on the CE routers
CE***#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, 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.x*.66 to network 0.0.0.0

     10.0.0.0/8 is variably subnetted, 8 subnets, 2 masks
S       10.1.0.0/24 is directly connected, Null0
       **** output omitted *****
S*   0.0.0.0/0 is directly connected, Serial0/0.114
```

- Use an **extended ping** command to verify that host addresses with the customer network can reach the Internet.

```
CEx1A#ping
Protocol [ip]:
Target IP address: 201.202.26.1
Repeat count [5]:
Datagram size [100]:
Timeout in seconds [2]:
Extended commands [n]: y
Source address or interface: 10.x.x1.49
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 201.202.26.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 100/135/193 ms
```

Lab 7-1 Answer Key: Separate Interface for Internet Connectivity

When you complete this activity, your router will be similar to the following, with differences that are specific to your pod.

Task 1: Establishing CE-PE Connectivity for Internet Access

Configuration steps on CEx1A:

```
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 CEx2B:

```
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 PEx1:

```
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 PEx2:

```
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 2: Establishing Routing Between the Customer and the Internet

Configuration steps on PE*x*1:

```
PEx1(config)#ip route 10.1.x1.0 255.255.255.0 150.x.x1.66
PEx1(config)#ip route 10.1.x2.0 255.255.255.0 150.x.x1.66
PEx1(config)#router bgp 65001
PEx1(config-router)#redistribute static
```

Configuration steps on PE*x*2:

```
PEx2(config)#ip route 10.2.x1.0 255.255.255.0 150.x.x2.66
PEx2(config)#ip route 10.2.x2.0 255.255.255.0 150.x.x2.66
PEx2(config)#router bgp 65001
PEx2(config-router)#redistribute static
```

Configuration steps on CEx1A:

```
CEx1A(config)#ip route 0.0.0.0 0.0.0.0 serial0/0.114
CEx1A(config)router bgp 650x1
CEx1A(config-router)#network 0.0.0.0
```

Configuration steps on CEx2B:

```
CEx2B(config)#ip route 0.0.0.0 0.0.0.0 serial0/0.114
CEx2B(config)router bgp 650x2
CEx2B(config-router)#network 0.0.0.0
```

Lab 7-2: Multisite Internet Access

To provide optimum routing, the service provider has convinced the customer to provide Internet access to each site. Because of the multisite access, the routing will have to be converted from static to BGP.

Note	This conversion will require additional firewall and NAT services that are not addressed by this lab activity.

Activity Objective

In this activity, you will migrate customers to direct BGP Interface access. After completing this activity, you will be able to meet these objectives:

- Establish remote site CE-PE connectivity for Internet access
- Establish remote site routing between the customer and the Internet

Visual Objective

You will configure additional virtual links (emphasized in the figure here) between the routers (CE*x*1B and CE*x*2A) and their PE routers. You will put these circuits and those created in the previous lab in the global routing table. You will also configure a global BGP session between PE routers and CE routers to exchange Internet routes between the service provider and the customer.

The figure illustrates what you will accomplish in this activity.

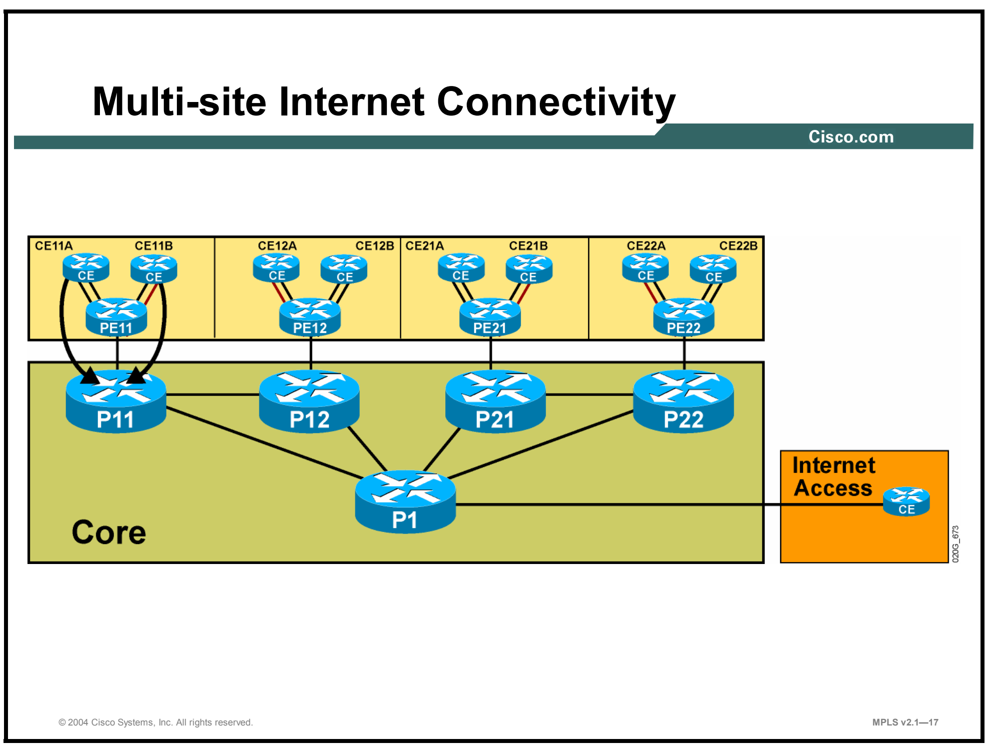

Required Resources

This is the resource required to complete this activity:

- Cisco IOS documentation

Command List

The table describes the commands used in this activity

Multisite Internet Access Command

Command	Description
ip route *prefix mask* **null 0**	Creates a summary route in the IP routing table.

Task 1: Establishing CE-PE Connectivity for Internet Access

Your service provider has already created a VPN to carry Internet traffic. You will need to join this VPN.

Activity Procedure

Complete these steps:

Step 1 Create a separate subinterface (S0/0.115) on the remaining router of the customer using the address information from this table.

Router ID	IP Address	DLCI
CEx1B	150.x.x1.130/28	115
CEx2A	150.x.x2.130/28	115

Step 2 Create a separate subinterface (0/0.115) on PE routers using the address information in this table.

Router ID	IP Address	DLCI
PEx1	150.x.x1.129/28	115
PEx2	150.x.x2.129/28	115

Activity Verification

You have completed this task when you attain these results:

- You have verified the status of the interface.

```
CEx1B#sh ip int s0/0.115
Serial0/0.115 is up, line protocol is up
  Internet address is 150.x.x1.130/28
  Broadcast address is 255.255.255.255
  Address determined by setup command
  MTU is 1500 bytes
  **** output omitted *******

PEx1#sh ip int s0/0.115
Serial0/0.115 is up, line protocol is up
  Internet address is 150.x.x1.129/28
  Broadcast address is 255.255.255.255
  Address determined by setup command
  MTU is 1500 bytes
  **** output omitted *******
```

Task 2: Establishing Routing Between the Customer and the Internet

The next task is to convert the interface created in Lab 7-1: Separate Interface for Internet Connectivity from static routing to an EBGP session. You then need to enable an EBGP session on the new interface.

Activity Procedure

Complete these steps:

Step 1 On your assigned central CE router (CE*x*1A or CE*x*2B), remove the network statement and passive interface command related to the WAN interface from the customer IGP process.

Step 2 Remove the network statement that refers to network 0.0.0.0 from BGP.

Step 3 Remove the 0.0.0.0 static route.

Step 4 Add the associated PE router as a BGP neighbor.

Step 5 On the associated PE router, add the associated CE router as a BGP neighbor.

Step 6 On your assigned CE router (CE*x*2A or CE*x*1B), add the associated PE router as a BGP neighbor.

Step 7 On the associated PE router, add the associated CE router as a BGP neighbor.

Activity Verification

You have completed this task when you attain these results:

- You have verified the status of the BGP neighbors.

```
PEx1#sh ip bgp sum
BGP router identifier 192.168.1.17, local AS number 65001
BGP table version is 41, main routing table version 41
36 network entries using 3636 bytes of memory
36 path entries using 1728 bytes of memory
18 BGP path attribute entries using 1080 bytes of memory
2 BGP rrinfo entries using 48 bytes of memory
3 BGP AS-PATH entries using 72 bytes of memory
9 BGP extended community entries using 320 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 6884 total bytes of memory
BGP activity 96/9 prefixes, 107/13 paths, scan interval 60 secs

Neighbor        V    AS MsgRcvd MsgSent  TblVer  InQ OutQ Up/Down  State/PfxRcd
150.x.x1.66     4 650xx      10       7      41    0    0 00:01:57        4
150.x.x1.130    4 650xx      10       7      41    0    0 00:01:57        4
192.168.100.129 4 65001     129     114      41    0    0 01:30:45       32

CEx1A#ping
Protocol [ip]:
Target IP address: 201.202.26.1
Repeat count [5]:
Datagram size [100]:
Timeout in seconds [2]:
Extended commands [n]: y
Source address or interface: 10.x.x1.49
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 201.202.26.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 100/135/193 ms
```

Lab 7-2 Answer Key: Multisite Internet Access

When you complete this activity, your router will be similar to the following, with differences that are specific to your pod.

Task 1: Establishing CE-PE Connectivity for Internet Access

Configuration steps on CE routers:

```
CEx**(config)#interface serial0/0.115 point-to-point
CEx**(config-subif)#ip add 150.x.x*.130 255.255.255.240
CEx**(config-subif)#frame-relay interface-dlci 115
```

Configuration steps on PE routers:

```
PEx*(config)#interface serial0/0.115 point-to-point
PEx*(config-subif)#ip add 150.x.x*.129 255.255.255.240
PEx*(config-subif)#frame-relay interface-dlci 115
```

Task 2: Establishing Routing Between the Customer and the Internet

Customer A

Configuration steps on CE*x*1A router:

```
CEx1A(config)#router ospf 1
CEx1A(config-router)#no passive-interface serial0/0.114
CEx1A(config-router)#no network 150.x.0.0 0.0.255.255 area 0
CEx1A(config-router)#router bgp 650x1
CEx1A(config-router)#no network 0.0.0.0
CEx1A(config-router)#neighbor 150.x.x1.65 remote 65001
CEx1A(config-router)#no ip route 0.0.0.0 0.0.0.0 Serial0/0.114
```

Configuration steps on PE*x*1 routers:

```
PEx1(config)#no ip route 10.1.x1.0 255.255.255.0 150.x.x1.66
PEx1(config)#no ip route 10.1.x2.0 255.255.255.0 150.x.x1.66
PEx1(config)#router bgp 65001
PEx1(config-router)#neighbor 150.x.x1.66 remote 650x1
```

Configuration steps on CE*x*2A router:

```
CEx2A(config)#router bgp 650x1
CEx2A(config-router)#neighbor 150.x.x2.129 remote 65001
```

Configuration steps on PE*x*2 routers:

```
PEx2(config)#router bgp 65001
PEx2(config-router)#neighbor 150.x.x2.130 remote 650x1
```


Customer B

Configuration steps on CE*x*2B router:

```
CEx2B(config)#router ospf 2
CEx2B(config-router)#no passive-interface Serial0/0.114
CEx2B(config-router)#no network 150.x.0.0 0.0.255.255 area 0
CEx2B(config-router)#router bgp 650x2
CEx2B(config-router)#no network 0.0.0.0
CEx2B(config-router)#neighbor 150.x.x2.65 remote 65001
CEx2B(config-router)#no ip route 0.0.0.0 0.0.0.0 Serial0/0.114
```

Configuration steps on PE*x*2 routers:

```
PEx2(config)#no ip route 10.2.x1.0 255.255.255.0 150.x.x2.66
PEx2(config)#no ip route 10.2.x2.0 255.255.255.0 150.x.x2.66
PEx2(config)#router bgp 65001
PEx2(config-router)#neighbor 150.x.x2.66 remote 650x2
```

Configuration steps on CE*x*1B router:

```
CEx1B(config)#router bgp 650x2
CEx1B(config-router)#neighbor 150.x.x1.129 remote 65001
```

Configuration steps on PE*x*1 routers:

```
PEx1(config)#router bgp 65001
PEx1(config-router)#neighbor 150.x.x1.130 remote 650x2
```

Lab 7-3: Internet Connectivity in an MPLS VPN

Internet connectivity in MPLS VPN-based networks can be achieved through a dedicated Internet VPN. The dedicated Internet VPN approach gives you 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 transport full Internet routing in an Internet VPN.

Activity Objective

In this activity, you will migrate the customer to a VPN for Internet access. After completing this activity, you will be able to meet these objectives:

- Establish central site CE-PE connectivity for Internet access
- Establish remote site CE-PE connectivity for Internet access

Visual Objective

In this activity, you will create a VPN (VRF) that will carry all Internet traffic, and then you will create connectivity between that VPN and the customer site. Each workgroup will be responsible for performing the configuration tasks on its PE router.

The figure illustrates what you will accomplish in this activity.

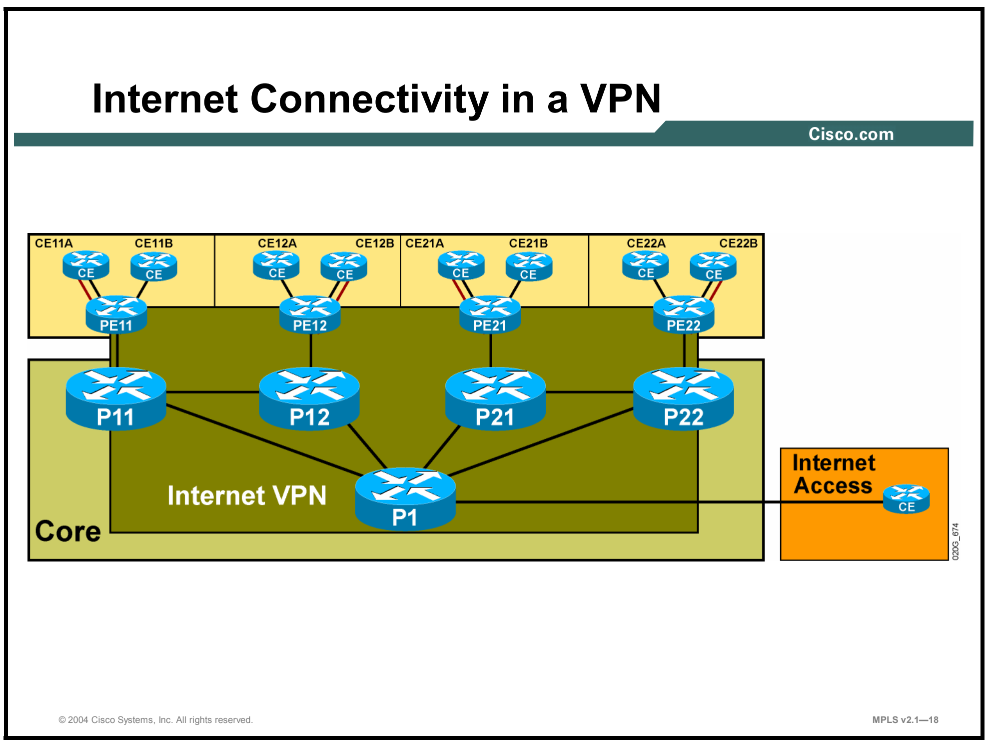

Required Resources

This is the resource required to complete this activity:

- Cisco IOS documentation

Command List

All commands used in this lab have been used in previous labs.

Task 1: Establishing Central Site Connectivity for Internet Access

Your service provider has already created a VPN to carry the Internet traffic. You will need to join this VPN.

Activity Procedure

Complete these steps:

Step 1 On your assigned PE router (PE*x*1 or PE*x*2), create a new Internet VPN VRF. The service provider has assigned an RT of 100:600 and a route distinguisher (RD) of 100:600 for all Internet-related VRFs.

Step 2 Place the interface (114) that is supporting the central site CE router (CE*x*1A or CE*x*2B) into the VRF.

Step 3 Remove the central site router neighbor statement from the unicast (global) address family.

Step 4 Add the central site router neighbor statement to the IPv4 VRF address family for the Internet VRF.

Activity Verification

You have completed this task when you attain these results:

- You have verified that the Internet routes being received by the central site CE route are coming from its PE neighbor.

```
CEx1A#sh ip rou
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
B    201.202.20.0/24 [20/0] via 150.x.x1.65, 01:54:31
B    202.100.36.0/24 [20/0] via 150.x.x1.65, 01:54:31
B    207.69.48.0/24 [20/0] via 150.x.x1.65, 01:54:31
B    201.202.21.0/24 [20/0] via 150.x.x1.65, 01:54:31
B    202.100.37.0/24 [20/0] via 150.x.x1.65, 01:54:31
B    207.69.49.0/24 [20/0] via 150.x.x1.65, 01:54:31
B    201.202.22.0/24 [20/0] via 150.x.x1.65, 01:54:31
B    202.100.38.0/24 [20/0] via 150.x.x1.65, 01:54:31
```

```
B    201.202.23.0/24 [20/0] via 150.x.x1.65, 01:54:31
B    202.100.39.0/24 [20/0] via 150.x.x1.65, 01:54:32
B    202.100.32.0/24 [20/0] via 150.x.x1.65, 01:54:32
****************** output omitted *************************
```

Task 2: Establishing Remote Site Connectivity for Internet Access

Your service provider has already created a VPN to carry Internet traffic. You will need to join this VPN.

Activity Procedure

Complete these steps:

Step 1 On your assigned PE router (PE*x*1 or PE*x*2) that supports your remote CE router (CE*x*2A or CE*x*1B), place the interface (115) into the VRF.

Step 2 Remove the remote site router neighbor statement for the unicast (global) address family.

Step 3 Add the remote site router neighbor statement to the IPv4 VRF address family for the Internet VRF.

Activity Verification

You have completed this task when you attain these results:

- You have verified that the Internet routes being received by the central site CE router are coming from its PE neighbor.

```
CEx2A#sh ip rou
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

B    201.202.20.0/24 [20/0] via 150.x.x2.129, 01:44:20
B    202.100.36.0/24 [20/0] via 150.x.x2.129, 01:44:20
B    207.69.48.0/24 [20/0] via 150.x.x2.129, 01:44:20
B    201.202.21.0/24 [20/0] via 150.x.x2.129, 01:44:20
B    202.100.37.0/24 [20/0] via 150.x.x2.129, 01:44:20
B    207.69.49.0/24 [20/0] via 150.x.x2.129, 01:44:20
B    201.202.22.0/24 [20/0] via 150.x.x2.129, 01:44:20
B    202.100.38.0/24 [20/0] via 150.x.x2.129, 01:44:20
B    201.202.23.0/24 [20/0] via 150.x.x2.129, 01:44:20
```



```
B    202.100.39.0/24 [20/0] via 150.x.x2.129, 01:44:21
B    202.100.32.0/24 [20/0] via 150.x.x2.129, 01:44:21
B    202.100.33.0/24 [20/0] via 150.x.x2.129, 01:44:21
```

Lab 7-3 Answer Key: Internet Connectivity in an MPLS VPN

When you complete this activity, your router will be similar to the following, with differences that are specific to your pod.

Task 1: Establishing Central Site Connectivity for Internet Access

Configuration steps on PE routers:

```
PEx1(config)#ip vrf Internet
PEx1(config-vrf)#route-target both 100:600
PEx1(config-vrf)#rd 100:600
PEx1(config)#interface serial0/0.114
PEx1(config-subif)#ip vrf forwarding Internet
% Interface Serial0/0.114 IP address 150.x.x1.65 removed due to enabling VRF
Internet
PEx1(config-subif)#ip add 150.x.x1.65 255.255.255.240
PEx1(config)#router bgp 65001
PEx1(config-router)#no neighbor 150.x.x1.66 remote-as 650x1
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

PEx2(config)#ip vrf Internet
PEx2(config-vrf)#route-target both 100:600
PEx2(config-vrf)#rd 100:600
PEx2(config)#interface serial0/0.114
PEx2(config-subif)#ip vrf forwarding Internet
% Interface Serial0/0.114 IP address 150.x.x2.65 removed due to enabling VRF
Internet
PEx2(config-subif)#ip add 150.x.x2.65 255.255.255.240
PEx2(config)#router bgp 65001
PEx2(config-router)#no neighbor 150.x.x2.66 remote-as 650x2
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
```


Task 2: Establishing Remote Site CE-PE Connectivity for Internet Access

Configuration steps on PE routers:

```
PEx1(config-vrf)#interface serial0/0.115

PEx1(config-subif)#ip vrf forward Internet

% Interface Serial0/0.115 IP address 150.x.x1.129 removed due to enabling VRF
Internet

PEx1(config-subif)#ip add 150.x.x1.129 255.255.255.240

PEx1(config-subif)#router bgp 65001

PEx1(config-router)#no neighbor 150.x.x1.130

PEx1(config-router)#address-family ipv4 vrf Internet

PEx1(config-router-af)#neighbor 150.x.x1.130 remote 650x2

PEx1(config-router-af)#neighbor 150.x.x1.130 activate


PEx2(config-vrf)#interface serial0/0.115

PEx2(config-subif)#ip vrf forward Internet

% Interface Serial0/0.115 IP address 150.x.x2.129 removed due to enabling VRF
Internet

PEx2(config-subif)#ip add 150.x.x2.129 255.255.255.240

PEx2(config-subif)#router bgp 65001

PEx2(config-router)#no neighbor 150.x.x2.130

PEx2(config-router)#address-family ipv4 vrf Internet

PEx2(config-router-af)#neighbor 150.x.x2.130 remote 650x1

PEx2(config-router-af)#neighbor 150.x.x2.130 activate
```