SPEDGE

Implementing Cisco Service Provider Next-Generation Edge Network Services

Version 1.0

Lab Guide

Text Part Number: 97-3157-01

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **www.cisco.com/go/offices.**

Table of Contents

Lab Guide

Overview

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

Outline

This guide includes these activities:

- Job Aids
- Lab 2-1: Implement MPLS Layer 3 VPN Backbones
- Lab 2-2: Connect Customers to MPLS Layer 3 VPNs
- Lab 2-3: Connect Advanced Customers to MPLS Layer 3 VPNs
- Lab 3-1: Establish Overlapping and Common Services Layer 3 VPNs
- Lab 3-2: Establish Internet Connectivity with an MPLS Layer 3 VPN
- Lab 3-3: Implement CSC
- Lab 4-1: Implement Layer 2 VPN (VPWS and VPLS)
- Tear-Out Section

Job Aids

These job aids are available to help you complete lab activities.

Pod Access Information

The instructor will provide you with the team and pod numbers, as well as access information for the other team and pod. Write down the information in the table for future reference.

Parameter	Default Value	Value
Team number	z = 1 to 4	
Pod number	x = 1, 3, 5, 7 or y = 2, 4, 6, 8	
Remote lab SSH access IP address	128.107.245.9	
Remote lab SSH access username	instr	
Remote lab SSH access password	testMe	
Pod PE (Cisco IOS XR) router username	root	
Pod PE (Cisco IOS XR) router password	1ronMan	
Pod CE, SW, and PE privileged-level password	cisco	

Device Information

This lab topology consists of four (4) teams and eight (8) pods. Two students will work in one pod and two pods will work in one team. Each pod has one switch and two routers. Two pods share one additional switch. All teams share the same core routers (P1 and P2).

The CE routers in both pods are running Cisco IOS Software. The first pod within a team (pod 1, 3, 5, or 7) will work on the PE router that is running Cisco IOS XR Software. The second pod within the same team (pod 2, 4, 6, or 8) will work on the PE router that is running Cisco IOS XE Software.

Devices in the lab are connected with Fast Ethernet and Gigabit Ethernet connections, and two teams have a redundant POS connection, as shown in the topology in the figure.

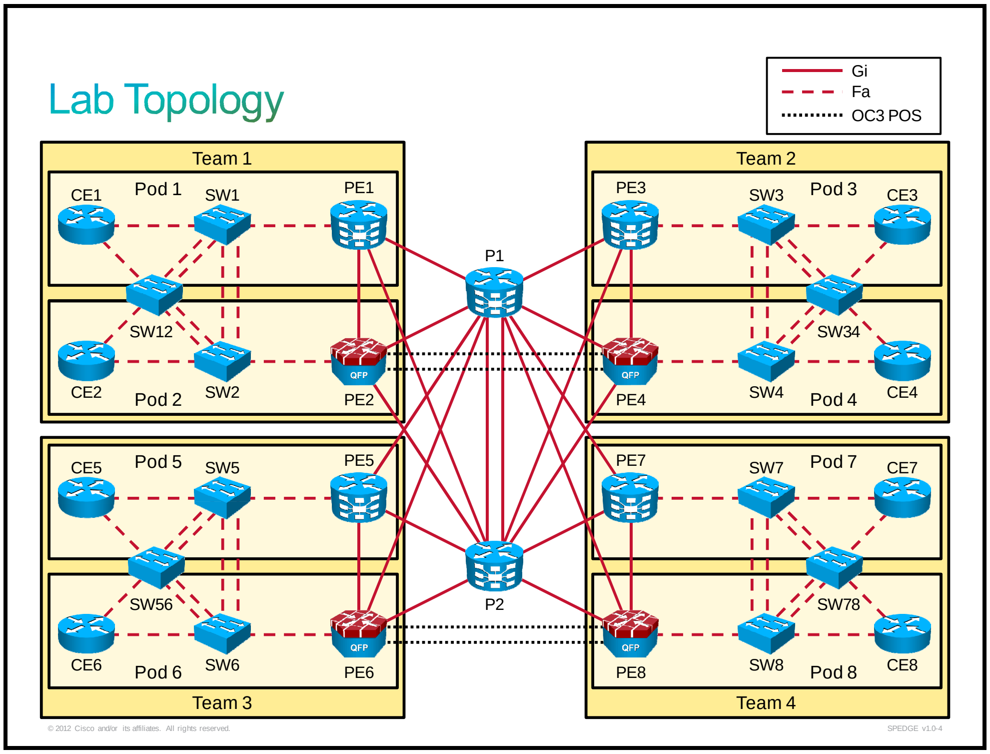

Device Roles and Loopback IP Addresses

Device Name	Device Role	Lo0 IPv4 Address	Lo0 IPv6 Address
CEx CEy	Cisco 2900 pod router	10.x.10.1/32 10.y.10.1/32	2001:db8:10:x:10::1/128 2001:db8:10:y:10::1/128
PEx PEy	Cisco ASR 9000 or Cisco ASR 1000 pod router	10.x.1.1/32 10.y.1.1/32	2001:db8:10:x:1::1/128 2001:db8:10:y:1::1/128
SWx SWy	Cisco ME340x pod switch	10.x.0.1/32 10.y.0.1/32	2001:db8:10:x:0::1/128 2001:db8:10:y:0::1/128
SWxy	Cisco ME340x pod switch, shared inside a team	10.xy.0.1/32	2001:db8:10:xy:0::1/128
P1	Cisco ASR 9000 core router	10.0.1.1/32	2001:db8:10:0:1::1/128
P2	Cisco ASR 9000 core router	10.0.2.1/32	2001:db8:10:0:2::1/128

The figure illustrates the interface identifications that are used in this lab setup.

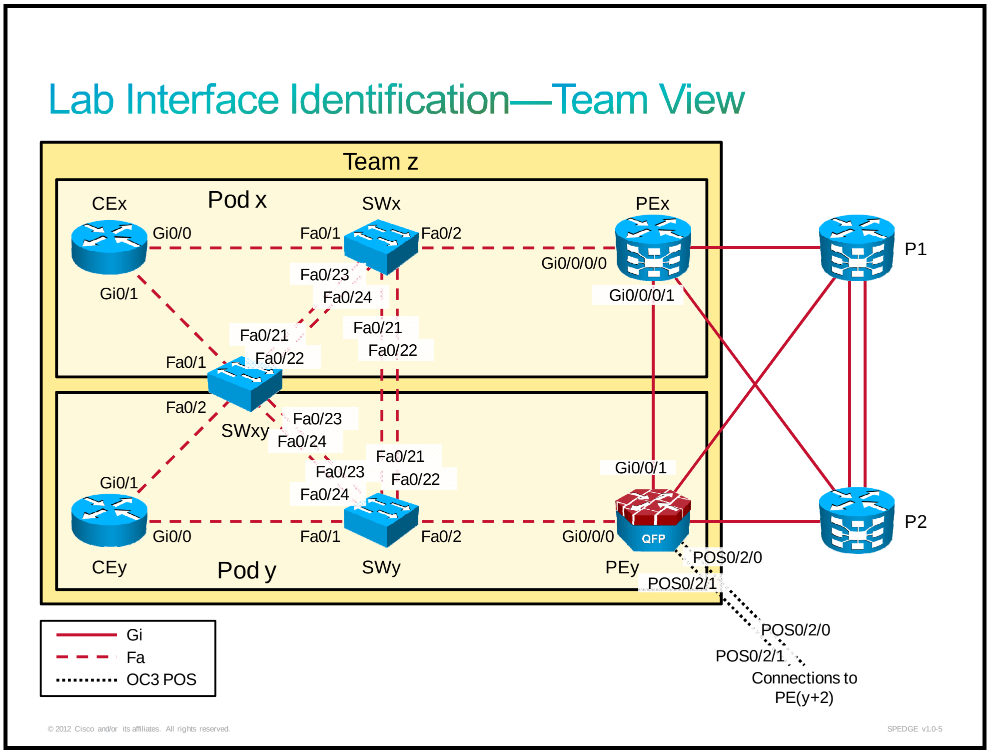

IP Addressing

The figure illustrates the IP addressing schemes that are used in this lab setup.

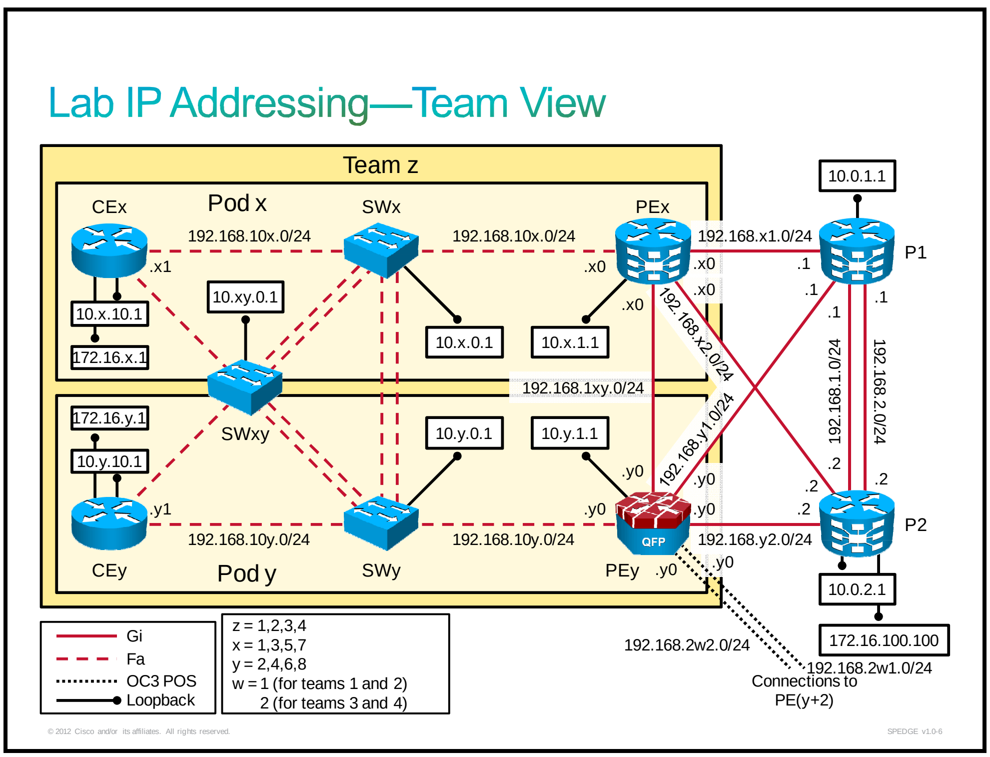

The following figure illustrates the management IP addresses that are used in this lab setup.

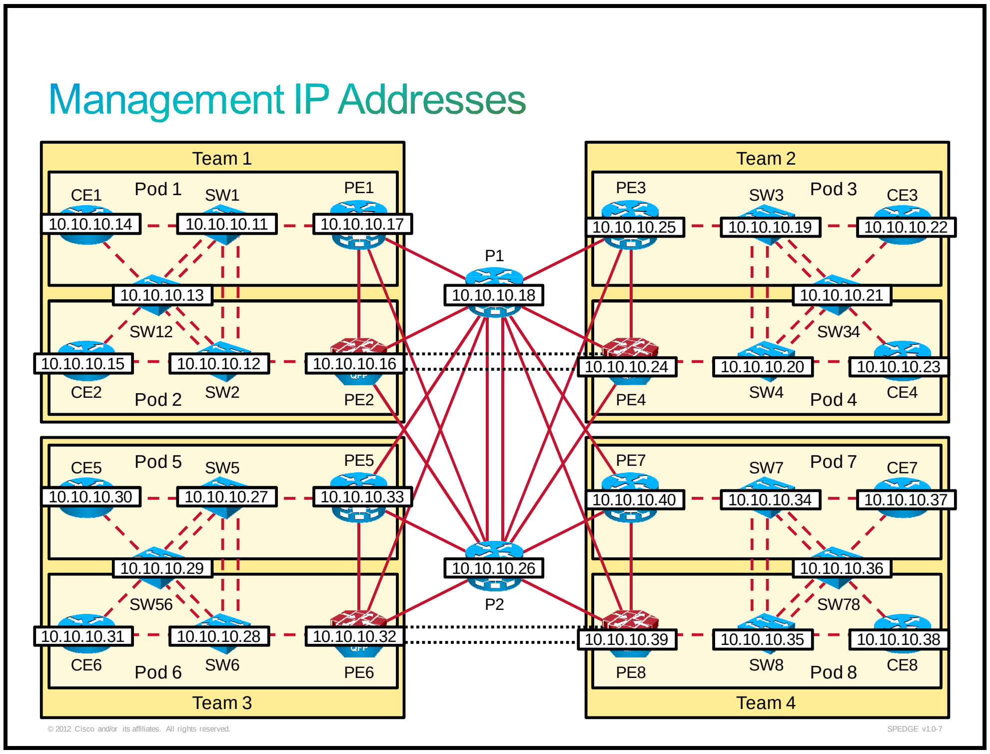

Note Replace the x or y with your pod number to get the IP addresses within your pod (x is for odd-numbered pods 1, 3, 5, and 7; y is for even-numbered pods 2, 4, 6, and 8). Replace the xy (where x < y) with the numbers of the pods within the same team (for example, 12, 34, 56, or 78) to get the IP addresses of the links between those pods.

Pod IP Addressing

Device	Interface	IPv4 Address	IPv6 Address
CEx	Gi0/0 Lo1	192.168.10x.x1/24 172.16.x.1/24	2001:db8:192:168:10x::x1/80 /
CEy	Gi0/0 Lo1	192.168.10y.y1/24 172.16.y.1/24	2001:db8:192:168:10y::y1/80 /
P1		192.168.x1.1/24	2001:db8:192:168:x1::1/80
		192.168.y1.1/24	2001:db8:192:168:y1::1/80
P2		192.168.x2.2/24	2001:db8:192:168:x2::2/80
		192.168.y2.2/24	2001:db8:192:168:y2::2/80
	Lo500	172.16.100.100	/
PE2	POS0/2/0	192.168.211.20/24	2001:db8:192:168:211::20/80
	POS0/2/1	192.168.212.20/24	2001:db8:192:168:212::20/80
PE4	POS0/2/0	192.168.211.40/24	2001:db8:192:168:211::40/80
	POS0/2/1	192.168.212.40/24	2001:db8:192:168:212::40/80
PE6	POS0/2/0	192.168.221.60/24	2001:db8:192:168:221::60/80
	POS0/2/1	192.168.222.60/24	2001:db8:192:168:222::60/80

Device	Interface	IPv4 Address	IPv6 Address
PE8	POS0/2/0	192.168.221.80/24	2001:db8:192:168:221::80/80
	POS0/2/1	192.168.222.80/24	2001:db8:192:168:222::80/80
PEx	Gi0/0/0/0	192.168.10x.x0/24	2001:db8:192:168:10x::x0/80
	Gi0/0/0/1	192.168.1xy.x0/24	2001:db8:192:168:1xy::x0/80
	Gi0/0/0/2	192.168.x1.x0/24	2001:db8:192:168:x1::x0/80
	Gi0/0/0/3	192.168.x2.x0/24	2001:db8:192:168:x2::x0/80
PEy	Gi0/0/0	192.168.10y.y0/24	2001:db8:192:168:10y::y0/80
	Gi0/0/1	192.168.1xy.y0/24	2001:db8:192:168:1xy::y0/80
	Gi0/0/2	192.168.y1.y0/24	2001:db8:192:168:y1::y0/80
	Gi0/0/3	192.168.y2.y0/24	2001:db8:192:168:y2::y0/80

Core IP Addressing

Device	Device IP Address	Peer	Peer IP Address
P1	192.168.1.1/24 2001:db8:192:168:1::1/80	P2	192.168.1.2/24 2001:db8:192:168:1::2/80
	192.168.2.1/24 2001:db8:192:168:2::1/80		192.168.2.2/24 2001:db8:192:168:2::2/80

Lab 2-1: Implement MPLS Layer 3 VPN Backbones

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

Activity Objective

In this lab activity, you will establish simple MPLS Layer 3 VPNs to support customer needs. Each pod is responsible for CE and PE router configurations that are related to the customer. This division of work between pods applies to all subsequent exercises in this course.

Note	Students from two pods will work in a team. The CE routers in both pods are running Cisco IOS Software. The first pod in the team will work on the PE router that is running Cisco IOS XR Software, and the second pod in the team will work on the PE router that is running Cisco IOS XE Software. Students in the same team should coordinate their lab activities.

You will work on different Cisco routers that are running Cisco IOS (c2900), Cisco IOS XE (asr1001), and Cisco IOS XR (asr9k) Software. After completing this activity, you will be able to meet these objectives:

- Enable LDP on your PE and P routers
- Configure the VRF tables that are necessary to support your customer
- Configure MP-BGP to establish routing between the PE routers

Visual Objective

The figure illustrates what you will accomplish in this activity.

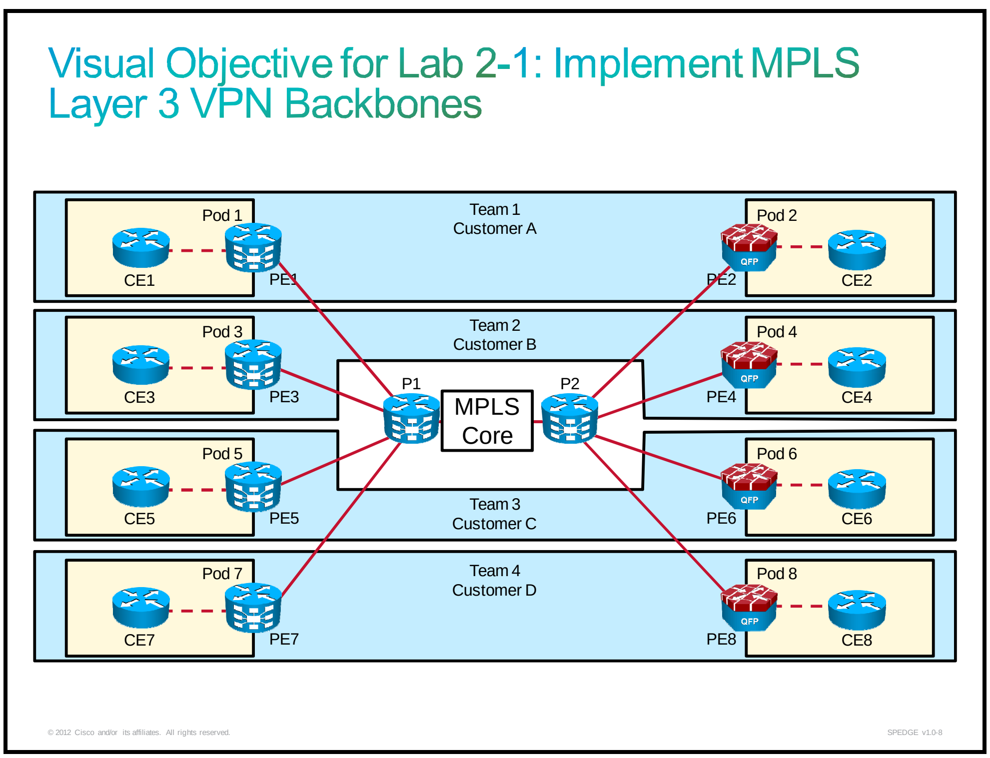

This activity contains tasks that enable you to configure your core MPLS VPN infrastructure.

VRF Assignments

This table gives the VRF names and RDs for the VRFs that are used in the lab.

VRF Details

Team	Description	VRF Name	VRF RD
Team 1	Customer A	Customer_1	1:210
Team 2	Customer B	Customer_2	1:220
Team 3	Customer C	Customer_3	1:230
Team 4	Customer D	Customer_4	1:240

Required Resources

These resources and equipment are required to complete this activity:

- A PC with access to the Internet
- An SSH client that is installed on the PC

Command List

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

Cisco IOS and IOS XE Commands

Command	Description
[no] shutdown	Enables or disables an interface
configure terminal	Enters configuration mode
mpls label protocol {ldp | tdp | both }	Specifies the label distribution protocol to be used on a given interface or globally
mpls ip	Enables MPLS forwarding of IPv4 packets along normally routed paths for the platform. The **mpls ip** command can be used in global configuration mode (for TE) but must be used in interface configuration mode for LDP to become active.
interface *interface*	Enters interface configuration mode
ping *dest_IP* **source** *source_IP*	Verifies connectivity between source IP and destination IP
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 activate	Activates an exchange of routes from the address family under the configuration for the specified neighbor
neighbor next-hop-self	Configures the router as the next hop for a BGP-speaking neighbor or peer group (in router configuration mode)
neighbor remote-as	Adds an entry to the BGP or MP-BGP neighbor table (in router configuration mode)
neighbor send-community	Specifies that a community attribute should be sent to a BGP neighbor (in address family or router configuration mode)
neighbor update-source	Has Cisco IOS Software allow IBGP sessions to use any operational interface for TCP connections (in router configuration mode)
router bgp *as-number*	Selects BGP configuration
route-target import | export *value*	Assigns an RT to a VRF
rd *value*	Assigns an RD to a VRF
address-family vpnv4	Selects VPNv4 address-family configuration
address-family ipv4	Selects IPv4 address-family configuration

Cisco IOS XR Commands

Command	Description
[no] shutdown	Enables or disables an interface
commit	Commits changes to the running configuration
configure terminal	Enters configuration mode
mpls ldp	Enters MPLS LDP configuration submode
router-id [*type number* | *ip-address*]	Specifies the router ID of the local node. In Cisco IOS XR, the router ID is specified as an interface name or IP address.
interface *type number*	Enters interface configuration mode for LDP (LDP mode)
interface *interface*	Enters interface configuration mode (in global configuration mode)
commit	Commits changes to the running configuration
configure terminal	Enters configuration mode
ipv4 | ipv6 address *ip_address/len*	Sets the IPv4 or IPv6 address for an interface and the subnet mask using the prefix length format
interface *interface*	Enters interface configuration mode (in global configuration mode)
ping *dest_IP* **source** *source_IP*	Verifies connectivity between source IP and destination IP (IPv4 and IPv6)
vrf *vrf-name*	Creates a VRF table
address-family ipv4 vrf *vrf-name*	Selects a per-VRF instance of a routing protocol
address-family vpnv4 unicast	Selects VPNv4 address-family configuration
address-family ipv4 unicast	Selects IPv4 address-family configuration
import | export route-target *value*	Assigns an RT to a VRF
neighbor ip-address activate	Activates an exchange of routes from the address family under the configuration for the specified neighbor
neighbor next-hop-self	Configures the router as the next hop for a BGP-speaking neighbor or peer group (in router configuration mode)
neighbor remote-as	Adds an entry to the BGP or MP-BGP neighbor table (in router configuration mode)
neighbor send-community	Specifies that a community attribute should be sent to a BGP neighbor (in address family or router configuration mode)
neighbor update-source	Has Cisco IOS XR software allow IBGP sessions to use any operational interface for TCP connections (in router configuration mode)
router bgp *as-number*	Selects BGP configuration

Task 1: Configure the VRF Tables Necessary to Support the Customer

In this task and the following tasks, you will first enable LDP on the PE routers and then establish simple VPNs for the customer. Each pod is responsible for all PE router configurations that are related to the customer. This division of work between pods applies to all subsequent exercises in this course. All P routers are preconfigured.

Activity Procedure

Complete these steps to prepare the configuration for the routers in your pod. You will work with students from other pods to finish this task.

Step 1 Enable LDP on the interface that is facing the P router.

Step 2 Create a VRF instance on the PE router. Use the “VRF Details” table for reference.

Step 3 Associate the PE-CE interface with the configured VRF. Use the details from the “VRF Details” table.

Activity Verification

Complete the verification of the lab activity:

- On each of your routers, verify that the interfaces to the P routers have been configured to use LDP.

```
RP/0/RSP0/CPU0:PEx# show mpls interface
Interface                  LDP      Tunnel   Enabled
-------------------------- -------- -------- --------
GigabitEthernet0/0/0/2     Yes      No       Yes

PEy#show mpls interface
Interface              IP            Tunnel   BGP Static Operational
GigabitEthernet0/0/3   Yes (ldp)     No       No  No     Yes
```

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

```
RP/0/RSP0/CPU0:PEx#show mpls ldp neighbor
Peer LDP Identifier: 10.0.1.1:0
  TCP connection: 10.0.1.1:646 - 10.3.1.1:43457
  Graceful Restart: Yes (Reconnect Timeout: 120 sec, Recovery: 0 sec)
  Session Holdtime: 180 sec
  State: Oper; Msgs sent/rcvd: 9891/9906; Downstream-Unsolicited
  Up time: 5d22h
  LDP Discovery Sources:
    GigabitEthernet0/0/0/2
    Targeted Hello (10.3.1.1 -> 10.0.1.1, active)
  Addresses bound to this peer:
    10.0.1.1         10.10.10.18      192.168.2.1      192.168.11.1
    192.168.31.1     192.168.51.1     192.168.61.1     192.168.71.1

PEy#show mpls ldp neighbor
    Peer LDP Ident: 10.0.2.1:0; Local LDP Ident 10.4.1.1:0
        TCP connection: 10.0.2.1.646 - 10.4.1.1.63621
        State: Oper; Msgs sent/rcvd: 7888/7893; Downstream
        Up time: 4d18h
        LDP discovery sources:
          GigabitEthernet0/0/3, Src IP addr: 192.168.42.2
        Addresses bound to peer LDP Ident:
          10.10.10.26     10.0.2.1        209.165.200.225 209.165.201.1
          209.165.202.129 192.168.2.2     192.168.82.2    192.168.52.2
          192.168.42.2    192.168.62.2    192.168.22.2
```

```
RP/0/RSP0/CPU0:PEx#show mpls ldp discovery
Local LDP Identifier: 10.3.1.1:0
Discovery Sources:
  Interfaces:
    GigabitEthernet0/0/0/2 : xmit/recv
      LDP Id: 10.0.1.1:0, Transport address: 10.0.1.1
          Hold time: 10 sec (local:10 sec, peer:10 sec)

  Targeted Hellos:
    10.3.1.1 -> 10.0.1.1 (active), xmit/recv
      LDP Id: 10.0.1.1:0
          Hold time: 90 sec (local:90 sec, peer:90 sec)

PEy#show mpls ldp discovery
 Local LDP Identifier:
    10.4.1.1:0
    Discovery Sources:
    Interfaces:
        GigabitEthernet0/0/3 (ldp): xmit/recv
            LDP Id: 10.0.2.1:0
```

- Verify that you have properly configured your VRF tables by using the **show ip vrf detail** command. Your output should be similar to this example:

```
RP/0/RSP0/CPU0:PEX#show vrf all detail
VRF Customer_1; RD not set; VPN ID not set
Description not set
Interfaces:
  GigabitEthernet0/0/0/0
Address family IPV4 Unicast
  Import VPN route-target communities:
    RT:1:210
  Export VPN route-target communities:
    RT:1:210
  No import route policy
  No export route policy
Address family IPV6 Unicast
  No import VPN route-target communities
  No export VPN route-target communities
  No import route policy
  No export route policy

PEY#show ip vrf detail
VRF Customer_1 (VRF Id = 2); default RD 1:210; default VPNID <not set>
  Interfaces:
    Gi0/0/0
VRF Table ID = 2
  Export VPN route-target communities
    RT:1:210
  Import VPN route-target communities
    RT:1:210
  No import route-map
  No export route-map
  VRF label distribution protocol: not configured
  VRF label allocation mode: per-prefix
```


Task 2: Configure MP-BGP to Establish Routing Between the PE Routers

In this task, you will configure MP-BGP between the PE routers in your pod. You will configure an IBGP session with a route reflector with the IP address 10.0.1.1.

Pod x will configure MP-BGP on PEx (Cisco IOS XR Software), and Pod y will perform the same task on PEy (Cisco IOS XE Software).

Activity Procedure

Complete these steps:

Step 1 Activate the BGP process on your assigned router, using AS **64500** as the AS number. Configure an IBGP neighbor relationship with a route reflector router (**10.0.1.1**). Use Loopback0 as the source interface for the BGP session.

Step 2 Enable **vpnv4 unicast address-family** and activate the configured neighbor for that address family. Configure **next-hop-self** functionality. On the routers with the Cisco IOS XE operating system, configure the router to send standard and extended communities with route updates.

Step 3 Wait for the other pod to finish configuration and then run the verification steps.

Activity Verification

Complete the verification of the lab activity:

- Display the BGP neighbor information and ensure that BGP sessions have been established between the two PE routers.

```
RP/0/RSP0/CPU0:PE1#sh bgp vpnv4 unicast summary
BGP router identifier 10.1.1.1, local AS number 64500
BGP generic scan interval 60 secs
BGP table state: Active
Table ID: 0x0   RD version: 3889240856
BGP main routing table version 1
BGP scan interval 60 secs

BGP is operating in STANDALONE mode.


Process       RcvTblVer   bRIB/RIB   LabelVer  ImportVer  SendTblVer
StandbyVer
Speaker               1          1          1          1           1
1

Neighbor     Spk    AS MsgRcvd MsgSent   TblVer  InQ OutQ  Up/Down  St/PfxRcd
10.0.1.1       0 64500      28      25        1    0    0 00:22:39          0

RP/0/RSP0/CPU0:PE3#

PE2#sh bgp vpnv4 unicast all summary
BGP router identifier 10.2.1.1, local AS number 64500
BGP table version is 1, main routing table version 1

Neighbor    V       AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.0.1.1    4    64500      29      27        1    0    0 00:20:31        0
PE2#

RP/0/RSP0/CPU0:PE1#sh bgp neighbor
BGP neighbor is 10.0.1.1
 Remote AS 64500, local AS 64500, internal link
 Remote router ID 10.0.1.1
  BGP state = Established, up for 00:25:59
  Last read 00:00:54, Last read before reset 00:00:00
```

```
  Hold time is 180, keepalive interval is 60 seconds
  Configured hold time: 180, keepalive: 60, min acceptable hold time: 3
  Last write 00:00:54, attempted 19, written 19
  Second last write 00:01:54, attempted 19, written 19
  Last write before reset 00:00:00, attempted 0, written 0
  Second last write before reset 00:00:00, attempted 0, written 0
  Last write pulse rcvd  Dec  7 12:50:49.163 last full not set pulse count 56
  Last write pulse rcvd before reset 00:00:00
  Socket not armed for io, armed for read, armed for write
  Last write thread event before reset 00:00:00, second last 00:00:00
  Last KA expiry before reset 00:00:00, second last 00:00:00
  Last KA error before reset 00:00:00, KA not sent 00:00:00
  Last KA start before reset 00:00:00, second last 00:00:00
  Precedence: internet
  Neighbor capabilities:
    Route refresh: advertised and received
    4-byte AS: advertised and received
    Address family VPNv4 Unicast: advertised and received
  Received 31 messages, 0 notifications, 0 in queue
  Sent 28 messages, 0 notifications, 0 in queue
  Minimum time between advertisement runs is 0 secs

 For Address Family: VPNv4 Unicast
  BGP neighbor version 1
  Update group: 0.1 Filter-group: 0.3  No Refresh request being processed
  NEXT_HOP is always this router
  Route refresh request: received 0, sent 0
  0 accepted prefixes, 0 are bestpaths
  Cumulative no. of prefixes denied: 0.
  Prefix advertised 0, suppressed 0, withdrawn 0
  Maximum prefixes allowed 524288
  Threshold for warning message 75%, restart interval 0 min
  AIGP is enabled
  An EoR was received during read-only mode
  Last ack version 1, Last synced ack version 0
  Outstanding version objects: current 0, max 0
  Additional-paths operation: None

  Connections established 1; dropped 0
  Local host: 10.1.1.1, Local port: 24639
  Foreign host: 10.0.1.1, Foreign port: 179
  Last reset 00:00:00
RP/0/RSP0/CPU0:PE1#

PE2#sh ip bgp neighbors
BGP neighbor is 10.0.1.1,  remote AS 64500, internal link
  BGP version 4, remote router ID 10.0.1.1
  BGP state = Established, up for 00:24:51
  Last read 00:00:46, last write 00:00:22, hold time is 180, keepalive
interval is 60 seconds
  Neighbor sessions:
    1 active, is not multisession capable (disabled)
  Neighbor capabilities:
    Route refresh: advertised and received(new)
    Four-octets ASN Capability: advertised and received
    Address family IPv4 Unicast: advertised and received
    ipv4 MPLS Label capability: received
    Address family VPNv4 Unicast: advertised and received
    Multisession Capability:
  Message statistics:
    InQ depth is 0
    OutQ depth is 0

                         Sent       Rcvd
    Opens:                  1          1
    Notifications:          0          0
    Updates:                2          6
    Keepalives:            29         26
```

```
    Route Refresh:          0          0
    Total:                 32         33
  Default minimum time between advertisement runs is 0 seconds

 For address family: IPv4 Unicast
  Session: 10.0.1.1
  BGP table version 13, neighbor version 13/0
  Output queue size : 0
  Index 2, Advertise bit 0
  2 update-group member
  Slow-peer detection is disabled
  Slow-peer split-update-group dynamic is disabled
                               Sent       Rcvd
  Prefix activity:             ----       ----
    Prefixes Current:             0          4 (Consumes 352 bytes)
    Prefixes Total:               0          4
    Implicit Withdraw:            0          0
    Explicit Withdraw:            0          0
    Used as bestpath:           n/a          4
    Used as multipath:          n/a          0

                                 Outbound    Inbound
  Local Policy Denied Prefixes:    --------    -------
    Bestpath from this peer:              4        n/a
    Total:                                4          0
  Number of NLRIs in the update sent: max 0, min 0
  Last detected as dynamic slow peer: never
  Dynamic slow peer recovered: never
```

```
 For address family: VPNv4 Unicast
  Session: 10.0.1.1
  BGP table version 1, neighbor version 1/0
  Output queue size : 0
  Index 3, Advertise bit 0
  3 update-group member
  NEXT_HOP is always this router
  Slow-peer detection is disabled
  Slow-peer split-update-group dynamic is disabled
                               Sent       Rcvd
  Prefix activity:             ----       ----
    Prefixes Current:             0          0
    Prefixes Total:               0          0
    Implicit Withdraw:            0          0
    Explicit Withdraw:            0          0
    Used as bestpath:           n/a          0
    Used as multipath:          n/a          0

                                 Outbound    Inbound
  Local Policy Denied Prefixes:    --------    -------
    Total:                                0          0
  Number of NLRIs in the update sent: max 0, min 0
  Last detected as dynamic slow peer: never
  Dynamic slow peer recovered: never

  Address tracking is enabled, the RIB does have a route to 10.0.1.1
  Connections established 2; dropped 1
  Last reset 00:25:00, due to Peer closed the session
  Transport(tcp) path-mtu-discovery is enabled
  Graceful-Restart is disabled
Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Connection is ECN Disabled
Mininum incoming TTL 0, Outgoing TTL 255
Local host: 10.2.1.1, Local port: 41246
Foreign host: 10.0.1.1, Foreign port: 179
Connection tableid (VRF): 0

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

```
Event Timers (current time is 0x113D331):
Timer          Starts    Wakeups            Next
Retrans            31          1             0x0
TimeWait            0          0             0x0
AckHold            27         26             0x0
SendWnd             0          0             0x0
KeepAlive           0          0             0x0
GiveUp              0          0             0x0
PmtuAger            1          1             0x0
DeadWait            0          0             0x0
Linger              0          0             0x0

iss: 3184753686  snduna: 3184754351  sndnxt: 3184754351     sndwnd:  32165
irs: 3575395302  rcvnxt: 3575396431  rcvwnd:      15256  delrcvwnd:   1128

SRTT: 294 ms, RTTO: 346 ms, RTV: 52 ms, KRTT: 0 ms
minRTT: 1 ms, maxRTT: 300 ms, ACK hold: 200 ms
Status Flags: none
Option Flags: higher precendence, nagle, path mtu capable

Datagrams (max data segment is 1240 bytes):
Rcvd: 57 (out of order: 0), with data: 27, total data bytes: 1128
Sent: 60 (retransmit: 1 fastretransmit: 0),with data: 32, total data bytes:
664
PE2#
```


Lab 2-2: Connect Customers to MPLS Layer 3 VPNs

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

Activity Objective

In this activity, you will deploy various routing protocols as the PE-CE routing protocol in the VPN of your customer.

Note Students from two pods will work in a team. The CE routers in both pods are running Cisco IOS Software. The first pod in the team will work on the PE router that is running Cisco IOS XR Software, and the second pod in the team will work on the PE router that is running Cisco IOS XE Software. Students in the same team should coordinate their lab activities.

You will work on different Cisco routers that are running Cisco IOS (c2900), Cisco IOS XE (asr1001), and Cisco IOS XR (asr9k) Software. After completing this activity, you will be able to meet these objectives:

- Establish VPN routing using static routes between the PE and CE routers
- Establish VPN routing using RIP as the PE-CE routing protocol
- Establish VPN routing using EIGRP as the PE-CE routing protocol

Visual Objective

The figure illustrates what you will accomplish in this activity.

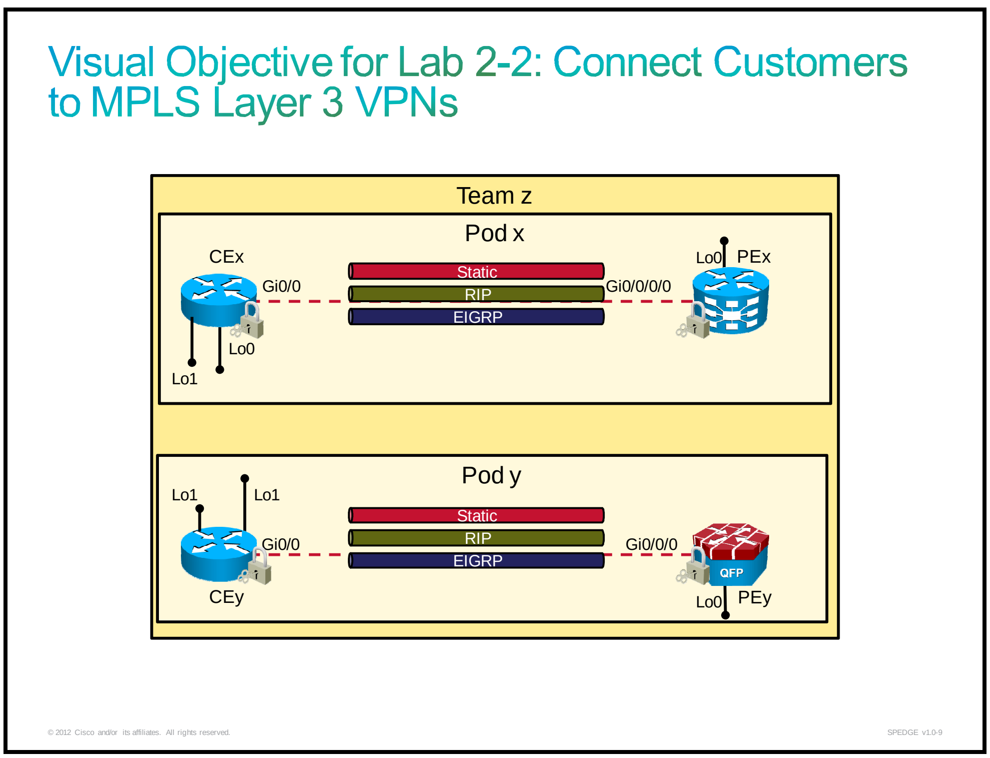

This activity contains tasks that enable you to configure a simple any-to-any VPN service for a customer.

You will test various simple PE-CE routing protocols between the PE and the CE routers.

VRF Assignments

This table gives the VRF RDs for VRFs that are used in this lab.

VRF Details

Team	Description	VRF RD
Team 1	Customer A	1:210
Team 2	Customer B	1:220
Team 3	Customer C	1:230
Team 4	Customer D	1:240

Required Resources

These resources and equipment are required to complete this activity:

- A PC with access to the Internet
- An SSH client that is installed on the PC

Command List

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

Cisco IOS and IOS XE Commands

Command	Description
[no] shutdown	Enables or disables an interface
configure terminal	Enters configuration mode
interface *interface*	Enters interface configuration mode
ping *dest_IP* **source** *source_IP*	Verifies connectivity between source IP and destination IP
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 activate	Activates an exchange of routes from the address family under the configuration for the specified neighbor
neighbor next-hop-self	Configures the router as the next hop for a BGP-speaking neighbor or peer group (in router configuration mode)
neighbor remote-as	Adds an entry to the BGP or MP-BGP neighbor table (in router configuration mode).
neighbor send-community	Specifies that a community attribute should be sent to a BGP neighbor (in address family or router configuration mode)
neighbor update-source	Has Cisco IOS software allow IBGP sessions to use any operational interface for TCP connections (in router configuration mode)
router bgp *as-number*	Selects BGP configuration
route-target import | export *value*	Assigns an RT to a VRF
rd *value*	Assigns an RD to a VRF
address-family vpnv4	Selects VPNv4 address-family configuration
address-family ipv4	Selects IPv4 address-family configuration
interface *interface*	Enters interface configuration mode
ping *dest_IP* **source** *source_IP*	Verifies connectivity between source IP and destination IP
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 activate	Activates an exchange of routes from the address family under the configuration for the specified neighbor
neighbor next-hop-self	Configures the router as the next hop for a BGP-speaking neighbor or peer group (in router configuration mode)

Command	Description
neighbor remote-as	Adds an entry to the BGP or MP-BGP neighbor table (in router configuration mode)
show running-config	Displays the running configuration
neighbor send-community	Specifies that a community attribute should be sent to a BGP neighbor (in address family or router configuration mode)
neighbor update-source	Has Cisco IOS software allow IBGP sessions to use any operational interface for TCP connections (in router configuration mode)
router bgp *as-number*	Selects BGP configuration
route-target import | export *value*	Assigns an RT to a VRF
rd *value*	Assigns an RD to a VRF
address-family vpnv4	Selects VPNv4 address-family configuration
address-family ipv4	Selects IPv4 address-family configuration
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**]	Redistribute 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 DUAL states and to debug possible DUAL problems.
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
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

Cisco IOS XR Commands

Command	Description
[no] shutdown	Enables or disables an interface
commit	Commits changes to the running configuration
configure terminal	Enters configuration mode
mpls ldp	Enters MPLS LDP configuration submode
router-id [*type number* | *ip-address*]	Specifies the router ID of the local node. In Cisco IOS XR Software, the router ID is specified as an interface name or IP address.
interface *type number*	Enters interface configuration mode for LDP (LDP mode)
interface *interface*	Enters interface configuration mode (in global configuration mode)
commit	Commits changes to the running configuration
configure terminal	Enters configuration mode
ipv4 | ipv6 address *ip_address/len*	Sets the IPv4 or IPv6 address for an interface and the subnet mask using the prefix length format
interface *interface*	Enters interface configuration mode (in global configuration mode)
ping *dest_IP* **source** *source_IP*	Verifies connectivity between source IP and destination IP (IPv4 and IPv6)
vrf *vrf-name*	Creates a VRF table
address-family ipv4 vrf *vrf-name*	Selects a per-VRF instance of a routing protocol
address-family vpnv4 unicast	Selects VPNv4 address-family configuration
address-family ipv4 unicast	Selects IPv4 address-family configuration
import | export route-target *value*	Assigns an RT to a VRF
neighbor ip-address activate	Activates an exchange of routes from the address family under the configuration for the specified neighbor
neighbor next-hop-self	Configures the router as the next hop for a BGP-speaking neighbor or peer group (in router configuration mode)
neighbor remote-as	Adds an entry to the BGP or MP-BGP neighbor table (in router configuration mode)
neighbor send-community	Specifies that a community attribute should be sent to a BGP neighbor (in address family or router configuration mode)
neighbor update-source	Has Cisco IOS XR software allow IBGP sessions to use any operational interface for TCP connections (in router configuration mode)
router bgp *as-number*	Selects BGP configuration
router eigrp *as-number*	Enters router configuration mode and creates an EIGRP routing process

Command	Description
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.
router ospf *process*	Starts an OSPF process
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
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

Task 1: Configuring Static Routes Between the PE and CE Routers

In this task, you will configure static routes between the PE and CE routers in your pod.

Activity Procedure

Complete these steps to prepare the configuration for the routers in your pod. You will work with students from other pods to finish this task.

Step 1 On the CE router, configure the loopback interface (Loopback1) with IP address **172.16.x.1/24** for Pod x and **172.16.y.1/24** for Pod y. The loopback interface will be used for customer network simulation.

Step 2 On the CE router, configure a default route with forwarding address **192.168.10x.x0** for Pod x and **192.168.10y.y0** for Pod y.

Step 3 On the PE router, configure a static route for the customer network. The interface that is facing the customer router is in VRF Customer_z, so you have to associate the static route with this VRF.

Step 4 In BGP processes, configure route redistribution under **ipv4 unicast address-family** (in VRF Customer_z). Redistribute the static and directly connected routes.

Activity Verification

Complete the verification of the lab activity:

- Verify the routing table on the CE router. A static default route should have been inserted into the routing table.

```
CE3#show ip route
<--- text omitted --->

Gateway of last resort is 192.168.103.30 to network 0.0.0.0

S*     0.0.0.0/0 [1/0] via 192.168.103.30
       10.0.0.0/32 is subnetted, 1 subnets
C         10.3.10.1 is directly connected, Loopback0
       172.16.0.0/16 is variably subnetted, 2 subnets, 2 masks
C         172.16.3.0/24 is directly connected, Loopback1
L         172.16.3.1/32 is directly connected, Loopback1
       192.168.103.0/24 is variably subnetted, 2 subnets, 2 masks
C         192.168.103.0/24 is directly connected, GigabitEthernet0/1
L         192.168.103.31/32 is directly connected, GigabitEthernet0/1
CE3#
```

- Verify the routing table on the PE router. Static and BGP routes should have been inserted into the VRF routing table.

```
RP/0/RSP0/CPU0:PE3#sh route vrf Customer_2
<--- text omitted --->

Gateway of last resort is not set

S    172.16.3.0/24 [1/0] via 192.168.103.31, 18:45:51
B    172.16.4.0/24 [200/0] via 10.4.1.1 (nexthop in vrf default), 01:02:38
C    192.168.103.0/24 is directly connected, 20:28:20, GigabitEthernet0/0/0/0
L    192.168.103.30/32 is directly connected, 20:28:20, GigabitEthernet0/0/0/0
B    192.168.104.0/24 [200/0] via 10.4.1.1 (nexthop in vrf default), 01:02:38
RP/0/RSP0/CPU0:PE3#
```

- Verify the advertised routes to the route reflector router.

```
RP/0/RSP0/CPU0:PE3#sh bgp vpnv4 unicast  neighbors 10.0.1.1 advertised-routes
```

```
Thu Dec  8 08:45:27.349 UTC
Network            Next Hop        From            AS Path
Route Distinguisher: 1:220
172.16.3.0/24      10.3.1.1        Local           ?
192.168.103.0/24   10.3.1.1        Local           ?

Processed 2 prefixes, 2 paths
RP/0/RSP0/CPU0:PE3#
```

- Verify connectivity between customer sites. Use the **ping** command on both CE routers.

```
CE3#ping 172.16.4.1 source loopback1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.4.1, timeout is 2 seconds:
Packet sent with a source address of 172.16.3.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
CE3#
CE4#ping 172.16.3.1 source loopback1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.3.1, timeout is 2 seconds:
Packet sent with a source address of 172.16.4.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
CE4#
```

- Trace the packet path between customer sites. Use the **traceroute** tool. You should see that different labels are assigned to IP packets.

```
CE3#traceroute 172.16.4.1 source loopback 1
Type escape sequence to abort.
Tracing the route to 172.16.4.1
VRF info: (vrf in name/id, vrf out name/id)
  1 192.168.103.30 0 msec 0 msec 0 msec
  2 192.168.31.1 [MPLS: Labels 16014/42 Exp 0] 0 msec 0 msec 0 msec
  3 192.168.1.2 [MPLS: Labels 16017/42 Exp 0] 4 msec 0 msec 0 msec
  4 192.168.104.40 [MPLS: Label 42 Exp 0] 0 msec 0 msec 0 msec
  5 192.168.104.41 0 msec 0 msec *
CE3#
CE4#trace 172.16.3.1 source loopback 1
Type escape sequence to abort.
Tracing the route to 172.16.3.1
VRF info: (vrf in name/id, vrf out name/id)
  1 192.168.104.40 0 msec 0 msec 0 msec
  2 192.168.42.2 [MPLS: Labels 16029/16026 Exp 0] 0 msec 0 msec 0 msec
  3 192.168.1.1 [MPLS: Labels 16004/16026 Exp 0] 4 msec 0 msec 0 msec
  4 192.168.31.30 [MPLS: Label 16026 Exp 0] 4 msec 0 msec 0 msec
  5 192.168.103.31 0 msec 0 msec *
CE4#
```

Task 2: Configure RIP as the PE-CE Routing Protocol

In this task, you will convert customer PE-CE routing from static to RIP.

Activity Procedure

Complete these steps:

Step 1 Remove static routes on the PE and CE routers.

Step 2 Configure RIP between the CE and PE routers. Advertise the customer network and the network of the segment between the PE and CE routers.

Step 3 Configure redistribution of RIP routes into BGP. Remove static route redistribution.

Step 4 Configure route redistribution of the customer BGP routes into RIP.

Activity Verification

Complete the verification of the lab activity:

- Verify the routing table on the CE router. RIP routes should have been inserted into the routing table.

```
CE3#show ip route
<--- text omitted --->

Gateway of last resort is not set

      10.0.0.0/32 is subnetted, 1 subnets
C        10.3.10.1 is directly connected, Loopback0
      172.16.0.0/16 is variably subnetted, 3 subnets, 2 masks
C        172.16.3.0/24 is directly connected, Loopback1
L        172.16.3.1/32 is directly connected, Loopback1
R        172.16.4.0/24
           [120/2] via 192.168.103.30, 00:00:24, GigabitEthernet0/1
      192.168.103.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.103.0/24 is directly connected, GigabitEthernet0/1
L        192.168.103.31/32 is directly connected, GigabitEthernet0/1
CE3#
```

- Verify the routing table on the PE router. RIP and BGP routes should have been inserted into the VRF routing table.

```
RP/0/RSP0/CPU0:PE3#sh route vrf Customer_2
<--- text omitted --->

Gateway of last resort is not set

R    172.16.3.0/24 [120/1] via 192.168.103.31, 00:56:22,
GigabitEthernet0/0/0/0
B    172.16.4.0/24 [200/1] via 10.4.1.1 (nexthop in vrf default), 00:36:23
C    192.168.103.0/24 is directly connected, 22:19:24, GigabitEthernet0/0/0/0
L    192.168.103.30/32 is directly connected, 22:19:24, GigabitEthernet0/0/0/0
B    192.168.104.0/24 [200/0] via 10.4.1.1 (nexthop in vrf default), 02:53:41
RP/0/RSP0/CPU0:PE3#
```

- Verify the advertised routes to the route reflector router.

```
RP/0/RSP0/CPU0:PE3#sh bgp vpnv4 unicast  neighbors 10.0.1.1 advertised-routes
Thu Dec  8 08:45:27.349 UTC
Network            Next Hop        From            AS Path
Route Distinguisher: 1:220
172.16.3.0/24      10.3.1.1        Local           ?
192.168.103.0/24   10.3.1.1        Local           ?

Processed 2 prefixes, 2 paths
RP/0/RSP0/CPU0:PE3#
```

- Verify the RIP database on the PE router.

```
RP/0/RSP0/CPU0:PE3#sh rip vrf Customer_2 database
Thu Dec  8 10:39:28.592 UTC

Routes held in RIP's topology database:
172.16.3.0/24
    [1] via 192.168.103.31, next hop 192.168.103.31, Uptime: 17s,
GigabitEthernet0/0/0/0
172.16.4.0/24
    [2] distance: 200    redistributed
172.16.0.0/16    auto-summary
192.168.103.0/24
    [0]    directly connected, GigabitEthernet0/0/0/0
#
```

- Verify connectivity between the customer sites. Use the **ping** command on both CE routers.

```
CE3#ping 172.16.4.1 source loopback1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.4.1, timeout is 2 seconds:
Packet sent with a source address of 172.16.3.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
CE3#
CE4#ping 172.16.3.1 source loopback1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.3.1, timeout is 2 seconds:
Packet sent with a source address of 172.16.4.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
CE4#
```

- Trace the packet path between customer sites. Use the **traceroute** tool. You should see that different labels are assigned to IP packets.

```
CE3#traceroute 172.16.4.1 source loopback 1
Type escape sequence to abort.
Tracing the route to 172.16.4.1
VRF info: (vrf in name/id, vrf out name/id)
  1 192.168.103.30 0 msec 0 msec 0 msec
  2 192.168.31.1 [MPLS: Labels 16014/42 Exp 0] 0 msec 0 msec 0 msec
  3 192.168.1.2 [MPLS: Labels 16017/42 Exp 0] 4 msec 0 msec 0 msec
  4 192.168.104.40 [MPLS: Label 42 Exp 0] 0 msec 0 msec 0 msec
  5 192.168.104.41 0 msec 0 msec *
CE3#
CE4#trace 172.16.3.1 source loopback 1
Type escape sequence to abort.
Tracing the route to 172.16.3.1
VRF info: (vrf in name/id, vrf out name/id)
  1 192.168.104.40 0 msec 0 msec 0 msec
  2 192.168.42.2 [MPLS: Labels 16029/16026 Exp 0] 0 msec 0 msec 0 msec
  3 192.168.1.1 [MPLS: Labels 16004/16026 Exp 0] 4 msec 0 msec 0 msec
  4 192.168.31.30 [MPLS: Label 16026 Exp 0] 4 msec 0 msec 0 msec
  5 192.168.103.31 0 msec 0 msec *
CE4#
```

Task 3: Configure EIGRP as the PE-CE Routing Protocol

In this activity, you will deploy EIGRP as the PE-CE routing protocol in the VPN of your customer.

Activity Procedure

Complete these steps:

Step 1 Remove the RIP configuration and configure EIGRP between the PE and CE routers. Use **1** for the EIGRP process number. Advertise the customer network and the network of the segment between the PE and CE routers.

Step 2 On your assigned PE router, configure redistribution of EIGRP routes into BGP. Remove redistribution of RIP routes.

Step 3 On your assigned PE router, configure redistribution of BGP routes into EIGRP. For the default metric, use these values:

- Bandwidth: **10000**
- Delay: **100**
- Reliability: **255**

- Loading: **1**
- MTU: **1500**

Activity Verification

Complete the verification of the lab activity:

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

```
RP/0/RSP0/CPU0:PE3#show eigrp vrf Customer_2 neighbors
Thu Dec  8 14:13:54.276 UTC

IPv4-EIGRP neighbors for AS(1) vrf Customer_2

H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
0   192.168.103.31          Gi0/0/0/0         14 00:38:02    4   200  0  3

RP/0/RSP0/CPU0:PE3#
```

- Verify the EIGRP topology database on the PE routers.

```
RP/0/RSP0/CPU0:PE3#show eigrp vrf Customer_z topology
IPv4-EIGRP Topology Table for AS(1)/ID(10.3.1.1) VRF: Customer_2

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

P 192.168.104.0/24, 1 successors, FD is 28160
        via VPNv4 Sourced (28160/0)
P 192.168.103.0/24, 1 successors, FD is 25856
        via Connected, GigabitEthernet0/0/0/0
P 172.16.4.0/24, 1 successors, FD is 156160
        via VPNv4 Sourced (156160/0)
P 172.16.3.0/24, 1 successors, FD is 153856
        via 192.168.103.31 (153856/128256), GigabitEthernet0/0/0/0

RP/0/RSP0/CPU0:PE3#
```

- Verify the routing table on the CE router. EIGRP routes should have been inserted into the routing table.

```
CE3#show ip route
<--- text omitted --->
Gateway of last resort is not set

      10.0.0.0/32 is subnetted, 1 subnets
C        10.3.10.1 is directly connected, Loopback0
      172.16.0.0/16 is variably subnetted, 3 subnets, 2 masks
C        172.16.3.0/24 is directly connected, Loopback1
L        172.16.3.1/32 is directly connected, Loopback1
D        172.16.4.0/24
           [90/158720] via 192.168.103.30, 00:32:05, GigabitEthernet0/1
      192.168.103.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.103.0/24 is directly connected, GigabitEthernet0/1
L        192.168.103.31/32 is directly connected, GigabitEthernet0/1
D     192.168.104.0/24
           [90/30720] via 192.168.103.30, 00:32:05, GigabitEthernet0/1
CE3#
```

- Verify the routing table on the PE router. EIGRP and BGP routes should have been inserted into the VRF routing table.

```
RP/0/RSP0/CPU0:PE3#sh route vrf Customer_2
<--- text omitted --->

Gateway of last resort is not set
```

```
D    172.16.3.0/24 [90/153856] via 192.168.103.31, 00:41:56,
GigabitEthernet0/0/0/0
B    172.16.4.0/24 [200/156160] via 10.4.1.1 (nexthop in vrf default),
00:35:16
C    192.168.103.0/24 is directly connected, 1d02h, GigabitEthernet0/0/0/0
L    192.168.103.30/32 is directly connected, 1d02h, GigabitEthernet0/0/0/0
B    192.168.104.0/24 [200/0] via 10.4.1.1 (nexthop in vrf default), 00:35:20
RP/0/RSP0/CPU0:PE3#
```

- Verify the advertised routes to the route reflector router.

```
RP/0/RSP0/CPU0:PE3#sh bgp vpnv4 unicast  neighbors 10.0.1.1 advertised-routes
Thu Dec  8 08:45:27.349 UTC
Network            Next Hop        From            AS Path
Route Distinguisher: 1:220
172.16.3.0/24      10.3.1.1        Local           ?
192.168.103.0/24   10.3.1.1        Local           ?

Processed 2 prefixes, 2 paths
RP/0/RSP0/CPU0:PE3#
```

- Verify connectivity between customer sites. Use the **ping** command on both CE routers.

```
CE3#ping 172.16.4.1 source loopback1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.4.1, timeout is 2 seconds:
Packet sent with a source address of 172.16.3.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
CE3#
CE4#ping 172.16.3.1 source loopback1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.3.1, timeout is 2 seconds:
Packet sent with a source address of 172.16.4.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
CE4#
```

- Trace the packet path between customer sites. Use the **traceroute** tool. You should see that different labels are assigned to IP packets.

```
CE3#traceroute 172.16.4.1 source loopback 1
Type escape sequence to abort.
Tracing the route to 172.16.4.1
VRF info: (vrf in name/id, vrf out name/id)
  1 192.168.103.30 0 msec 0 msec 0 msec
  2 192.168.31.1 [MPLS: Labels 16014/42 Exp 0] 0 msec 0 msec 0 msec
  3 192.168.1.2 [MPLS: Labels 16017/42 Exp 0] 4 msec 0 msec 0 msec
  4 192.168.104.40 [MPLS: Label 42 Exp 0] 0 msec 0 msec 0 msec
  5 192.168.104.41 0 msec 0 msec *
CE3#
CE4#trace 172.16.3.1 source loopback 1
Type escape sequence to abort.
Tracing the route to 172.16.3.1
VRF info: (vrf in name/id, vrf out name/id)
  1 192.168.104.40 0 msec 0 msec 0 msec
  2 192.168.42.2 [MPLS: Labels 16029/16026 Exp 0] 0 msec 0 msec 0 msec
  3 192.168.1.1 [MPLS: Labels 16004/16026 Exp 0] 4 msec 0 msec 0 msec
  4 192.168.31.30 [MPLS: Label 16026 Exp 0] 4 msec 0 msec 0 msec
  5 192.168.103.31 0 msec 0 msec *
CE4#
```

Lab 2-3: Connect Advanced Customers to MPLS Layer 3 VPNs

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

Activity Objective

In this activity, you will deploy various advanced routing protocols as the PE-CE routing protocol in the VPN of your customer.

Note	Students from two pods will work in a team. The CE routers in both pods are running Cisco IOS Software. The first pod in the team will work on the PE router that is running Cisco IOS XR Software, and the second pod in the team will work on the PE router that is running Cisco IOS XE Software. Students in the same team should coordinate their lab activities.

You will work on different Cisco routers that are running Cisco IOS (c2900), Cisco IOS XE (asr1001), and Cisco IOS XR (asr9k) Software. After completing this activity, you will be able to meet these objectives:

- Establish VPN routing using BGP as the PE-CE routing protocol
- Establish VPN routing using OSPF as the PE-CE routing protocol

Visual Objective

The figure illustrates what you will accomplish in this activity.

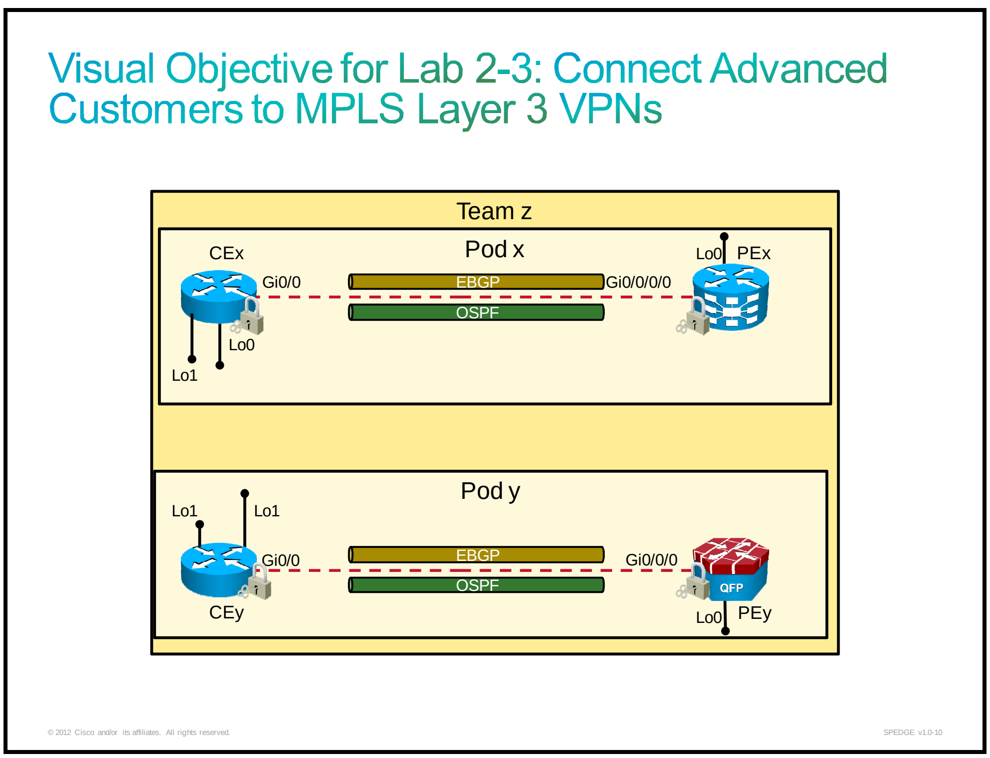

This activity contains tasks that enable you to configure a simple any-to-any VPN service for a customer.

You will test various PE-CE routing protocols between the PE and the CE routers.

VRF Assignments

This table gives the VRF RDs for VRFs that are used in this lab.

VRF Details

Team	Description	VRF RD
Team 1	Customer A	1:210
Team 2	Customer B	1:220
Team 3	Customer C	1:230
Team 4	Customer D	1:240

Required Resources

These resources and equipment are required to complete this activity:

- A PC with access to the Internet
- An SSH client that is installed on the PC

Command List

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

Cisco IOS and IOS XE Commands

Command	Description
[no] shutdown	Enables or disables an interface
configure terminal	Enters configuration mode
interface *interface*	Enters interface configuration mode
ping *dest_IP* **source** *source_IP*	Verifies connectivity between source IP and destination IP
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 activate	Activates an exchange of routes from the address family under the configuration for the specified neighbor
neighbor next-hop-self	Configures the router as the next hop for a BGP-speaking neighbor or peer group (in router configuration mode)
neighbor remote-as	Adds an entry to the BGP or MP-BGP neighbor table (in router configuration mode)
neighbor send-community	Specifies that a community attribute should be sent to a BGP neighbor (in address family or router configuration mode)
neighbor update-source	Has Cisco IOS software allow IBGP sessions to use any operational interface for TCP connections (in router configuration mode)
router bgp *as-number*	Selects BGP configuration
route-target import | export *value*	Assigns an RT to a VRF
rd *value*	Assigns an RD to a VRF
address-family vpnv4	Selects VPNv4 address-family configuration
address-family ipv4	Selects IPv4 address-family configuration
interface *interface*	Enters interface configuration mode
ping *dest_IP* **source** *source_IP*	Verifies connectivity between source IP and destination IP
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 activate	Activates an exchange of routes from the address family under the configuration for the specified neighbor
neighbor next-hop-self	Configures the router as the next hop for a BGP-speaking neighbor or peer group (in router configuration mode)

Command	Description
neighbor remote-as	Adds an entry to the BGP or MP-BGP neighbor table (in router configuration mode).
show running-config	Displays the running configuration
neighbor send-community	Specifies that a community attribute should be sent to a BGP neighbor (in address family or router configuration mode)
neighbor update-source	Has Cisco IOS software allow IBGP sessions to use any operational interface for TCP connections (in router configuration mode)
router bgp *as-number*	Selects BGP configuration
route-target import | export *value*	Assigns an RT to a VRF
rd *value*	Assigns an RD to a VRF
address-family vpnv4	Selects VPNv4 address-family configuration
address-family ipv4	Selects IPv4 address-family configuration
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 DUAL states and to debug possible DUAL problems.
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
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

Cisco IOS XR Commands

Command	Description
[no] shutdown	Enables or disables an interface
commit	Commits changes to the running configuration
configure terminal	Enters configuration mode
mpls ldp	Enters MPLS LDP configuration submode
router-id [*type number* | *ip-address*]	Specifies the router ID of the local node. In Cisco IOS XR, the router ID is specified as an interface name or IP address.
interface *type number*	Enters interface configuration mode for LDP (LDP mode)
interface *interface*	Enters interface configuration mode (in global configuration mode)
commit	Commits changes to the running configuration
configure terminal	Enters configuration mode
ipv4 | ipv6 address *ip_address/len*	Sets the IPv4 or IPv6 address for an interface and the subnet mask using the prefix length format
interface *interface*	Enters interface configuration mode (in global configuration mode)
ping *dest_IP* **source** *source_IP*	Verifies connectivity between source IP and destination IP (IPv4 and IPv6)
vrf *vrf-name*	Creates a VRF table
address-family ipv4 vrf *vrf-name*	Selects a per-VRF instance of a routing protocol
address-family vpnv4 unicast	Selects VPNv4 address-family configuration
address-family ipv4 unicast	Selects IPv4 address-family configuration
import | export route-target *value*	Assigns an RT to a VRF
neighbor ip-address activate	Activates an exchange of routes from the address family under the configuration for the specified neighbor
neighbor next-hop-self	Configures the router as the next hop for a BGP-speaking neighbor or peer group (in router configuration mode).
neighbor remote-as	Adds an entry to the BGP or MP-BGP neighbor table (in router configuration mode).
neighbor send-community	Specifies that a community attribute should be sent to a BGP neighbor (in address family or router configuration mode)
neighbor update-source	Has Cisco IOS XR software allow IBGP sessions to use any operational interface for TCP connections (in router configuration mode).
router bgp *as-number*	Selects BGP configuration
router eigrp *as-number*	Enters router configuration mode and creates an EIGRP routing process

Command	Description
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.
router ospf *process*	Starts an OSPF process
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
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

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

In this task, you will convert the CE-PE routing protocol of your customer from EIGRP to BGP.

Activity Procedure

Complete these steps:

Step 1 Remove the EIGRP configuration from all routers in your pod. Activate the BGP routing process on the CE routers in your pod. Use **6450x** for the AS number, where x is your pod number.

Step 2 Configure the **route-policy pass** command on the PEx router that will pass all routing updates.

Step 3 Activate the BGP neighbor relationship between the PE and CE routers in your pod. Use the CE-PE interface IP address as the source and destination IP addresses for the BGP session. Configure the **next-hop-self** command.

Step 4 Configure the CE router to advertise the customer network. On the PEx router (Cisco IOS XR Software) use the **route-policy pass** command to accept and send all routing updates to neighbors.

Activity Verification

Complete the verification of the lab activity:

- Verify the BGP table and BGP neighbor on the PE router.

```
RP/0/RSP0/CPU0:PE3#sh bgp vrf Customer_2 ipv4 unicast
Thu Dec  8 15:20:14.812 UTC
BGP VRF Customer_2, state: Active
BGP Route Distinguisher: 1:220
VRF ID: 0x60000012
BGP router identifier 10.3.1.1, local AS number 64500
BGP table state: Active
Table ID: 0xe0000012   RD version: 30
BGP main routing table version 30

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: 1:220 (default for vrf Customer_2)
*> 172.16.3.0/24      192.168.103.31           0             0 64503 i
*>i172.16.4.0/24      10.4.1.1            156160    100      0 ?
*> 192.168.103.0/24   0.0.0.0                  0         32768 ?
*>i192.168.104.0/24   10.4.1.1                 0    100      0 ?

Processed 4 prefixes, 4 paths
RP/0/RSP0/CPU0:PE3#

RP/0/RSP0/CPU0:PE3#sh bgp vrf Customer_2 ipv4 unicast summary
Thu Dec  8 15:20:23.473 UTC
BGP VRF Customer_2, state: Active
BGP Route Distinguisher: 1:220
VRF ID: 0x60000012
BGP router identifier 10.3.1.1, local AS number 64500
BGP table state: Active
Table ID: 0xe0000012   RD version: 30
BGP main routing table version 30

BGP is operating in STANDALONE mode.
```

```
Process      RcvTblVer   bRIB/RIB   LabelVer  ImportVer  SendTblVer  StandbyVer
Speaker             30         30         30         30          30          30

Neighbor       Spk    AS MsgRcvd MsgSent   TblVer  InQ OutQ  Up/Down  St/PfxRcd
192.168.103.31  0 64503      43      42       30    0    0 00:35:38          1

RP/0/RSP0/CPU0:PE3#
```

- Verify the routing table on the CE router. BGP routes should have been inserted into the routing table.

```
CE3#show ip route
<--- text omitted --->

      10.0.0.0/32 is subnetted, 1 subnets
C        10.3.10.1 is directly connected, Loopback0
      172.16.0.0/16 is variably subnetted, 3 subnets, 2 masks
C        172.16.3.0/24 is directly connected, Loopback1
L        172.16.3.1/32 is directly connected, Loopback1
B        172.16.4.0/24 [20/0] via 192.168.103.30, 00:20:14
      192.168.103.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.103.0/24 is directly connected, GigabitEthernet0/1
L        192.168.103.31/32 is directly connected, GigabitEthernet0/1
B     192.168.104.0/24 [20/0] via 192.168.103.30, 00:20:14
CE3#
```

- Verify the routing table on the PE router. BGP routes should have been inserted into the VRF routing table.

```
RP/0/RSP0/CPU0:PE3#sh route vrf Customer_2
<--- text omitted --->
Gateway of last resort is not set

B    172.16.3.0/24 [20/0] via 192.168.103.31, 00:21:19
B    172.16.4.0/24 [200/156160] via 10.4.1.1 (nexthop in vrf default),
01:41:44
C    192.168.103.0/24 is directly connected, 1d03h, GigabitEthernet0/0/0/0
L    192.168.103.30/32 is directly connected, 1d03h, GigabitEthernet0/0/0/0
B    192.168.104.0/24 [200/0] via 10.4.1.1 (nexthop in vrf default), 01:41:48
RP/0/RSP0/CPU0:PE3#
```

- Verify the advertised routes to the route reflector router.

```
RP/0/RSP0/CPU0:PE3#sh bgp vpnv4 unicast  neighbors 10.0.1.1 advertised-routes
Network            Next Hop        From            AS Path
Route Distinguisher: 1:220
172.16.3.0/24      10.3.1.1        192.168.103.31  64503i
192.168.103.0/24   10.3.1.1        Local           ?

Processed 2 prefixes, 2 paths
RP/0/RSP0/CPU0:PE3#
```

- Verify connectivity between customer sites. Use the **ping** command on both CE routers.

```
CE3#ping 172.16.4.1 source loopback1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.4.1, timeout is 2 seconds:
Packet sent with a source address of 172.16.3.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
CE3#
CE4#ping 172.16.3.1 source loopback1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.3.1, timeout is 2 seconds:
Packet sent with a source address of 172.16.4.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
CE4#
```

- Trace the packet path between customer sites. Use the **traceroute** tool. You should see that different labels are assigned to IP packets.

```
CE3#traceroute 172.16.4.1 source loopback 1
Type escape sequence to abort.
Tracing the route to 172.16.4.1
VRF info: (vrf in name/id, vrf out name/id)
  1 192.168.103.30 0 msec 0 msec 0 msec
  2 192.168.31.1 [MPLS: Labels 16014/42 Exp 0] 0 msec 0 msec 0 msec
  3 192.168.1.2 [MPLS: Labels 16017/42 Exp 0] 4 msec 0 msec 0 msec
  4 192.168.104.40 [MPLS: Label 42 Exp 0] 0 msec 0 msec 0 msec
  5 192.168.104.41 0 msec 0 msec *
CE3#
CE4#trace 172.16.3.1 source loopback 1
Type escape sequence to abort.
Tracing the route to 172.16.3.1
VRF info: (vrf in name/id, vrf out name/id)
  1 192.168.104.40 0 msec 0 msec 0 msec
  2 192.168.42.2 [MPLS: Labels 16029/16026 Exp 0] 0 msec 0 msec 0 msec
  3 192.168.1.1 [MPLS: Labels 16004/16026 Exp 0] 4 msec 0 msec 0 msec
  4 192.168.31.30 [MPLS: Label 16026 Exp 0] 4 msec 0 msec 0 msec
  5 192.168.103.31 0 msec 0 msec *
CE4#
```

Task 2: Configure OSPF as the PE-CE Routing Protocol

In this task, you will migrate the CE-PE routing protocol to OSPF. After completing this activity, you will be able to establish VPN routing using OSPF.

Activity Procedure

Complete these steps:

Step 1 Remove the EBGP configuration from the PE and CE routers. Configure OSPF between the CE and PE routers in your pod. Use an OSPF process ID of **2z0** (where z is your team number).

Step 2 On the CE router, configure the OSPF network type of the Loopback1 interface to point-to-point. Using the OSPF process, advertise the networks of Loopback1 and the PE-CE interfaces. All networks should be in OSPF Area 0.

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

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

Activity Verification

Complete the verification of the lab activity:

- Verify the OSPF database and all OSPF neighbors on the PE router.

```
RP/0/RSP0/CPU0:PE3#sh ospf vrf Customer_2 neighbor
Fri Dec  9 10:07:25.273 UTC

* Indicates MADJ interface

Neighbors for OSPF 220, VRF Customer_2

Neighbor ID     Pri   State           Dead Time   Address         Interface
172.16.3.1      1     FULL/DR         00:00:38    192.168.103.31
GigabitEthernet0/0/0/0
    Neighbor is up for 00:49:32
```

```
Total neighbor count: 1
RP/0/RSP0/CPU0:PE3#

RP/0/RSP0/CPU0:PE3#sh ospf vrf Customer_2 database
Fri Dec  9 10:07:51.225 UTC


            OSPF Router with ID (10.3.1.1) (Process ID 220, VRF Customer_2)

                Router Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum Link count
10.3.1.1        10.3.1.1        731         0x80000004 0x002e02 1
172.16.3.1      172.16.3.1      1123        0x80000007 0x005e9c 2

                Net Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum
192.168.103.31  172.16.3.1      1123        0x80000002 0x00a117

                Type-5 AS External Link States

Link ID         ADV Router      Age         Seq#       Checksum Tag
172.16.4.0      10.3.1.1        486         0x80000002 0x0031ee 3489725428
192.168.104.0   10.3.1.1        731         0x80000002 0x00aa65 3489725428
RP/0/RSP0/CPU0:PE3#
```

- Verify the OSPF database and all OSPF neighbors on the CE router.

```
CE3#sh ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
10.3.1.1          1   FULL/BDR        00:00:33    192.168.103.30
GigabitEthernet0/1
CE3#sh ip ospf dat
CE3#sh ip ospf database

            OSPF Router with ID (172.16.3.1) (Process ID 220)

                Router Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum Link count
10.3.1.1        10.3.1.1        723         0x80000004 0x002E02 1
172.16.3.1      172.16.3.1      1113        0x80000007 0x005E9C 2

                Net Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum
192.168.103.31  172.16.3.1      1113        0x80000002 0x00A117

                Type-5 AS External Link States

Link ID         ADV Router      Age         Seq#       Checksum Tag
172.16.4.0      10.3.1.1        478         0x80000002 0x0031EE 3489725428
192.168.104.0   10.3.1.1        723         0x80000002 0x00AA65 3489725428
CE3# #
```

- Verify the routing table on the CE router. BGP routes should have been inserted into the routing table.

```
CE3#show ip route
<--- text omitted --->
Gateway of last resort is not set

      10.0.0.0/32 is subnetted, 1 subnets
C        10.3.10.1 is directly connected, Loopback0
      172.16.0.0/16 is variably subnetted, 3 subnets, 2 masks
C        172.16.3.0/24 is directly connected, Loopback1
```

```
L        172.16.3.1/32 is directly connected, Loopback1
O E2     172.16.4.0/24
           [110/2] via 192.168.103.30, 00:43:49, GigabitEthernet0/1
      192.168.103.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.103.0/24 is directly connected, GigabitEthernet0/1
L        192.168.103.31/32 is directly connected, GigabitEthernet0/1
O E2  192.168.104.0/24
           [110/1] via 192.168.103.30, 00:46:13, GigabitEthernet0/1
CE3#
```

- Verify the routing table on the PE router. BGP routes should have been inserted into the VRF routing table.

```
RP/0/RSP0/CPU0:PE3#sh route vrf Customer_2
<--- text omitted --->
Gateway of last resort is not set

O    172.16.3.0/24 [110/2] via 192.168.103.31, 00:54:06,
GigabitEthernet0/0/0/0
B    172.16.4.0/24 [200/2] via 10.4.1.1 (nexthop in vrf default), 00:44:31
C    192.168.103.0/24 is directly connected, 1d21h, GigabitEthernet0/0/0/0
L    192.168.103.30/32 is directly connected, 1d21h, GigabitEthernet0/0/0/0
B    192.168.104.0/24 [200/0] via 10.4.1.1 (nexthop in vrf default), 00:46:00
RP/0/RSP0/CPU0:PE3##
```

- Verify connectivity between customer sites. Use the **ping** command on both CE routers.

```
CE3#ping 172.16.4.1 source loopback1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.4.1, timeout is 2 seconds:
Packet sent with a source address of 172.16.3.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
CE3#
CE4#ping 172.16.3.1 source loopback1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.3.1, timeout is 2 seconds:
Packet sent with a source address of 172.16.4.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
CE4#
```

- Trace the packet path between customer sites. Use the **traceroute** tool. You should see that different labels are assigned to IP packets.

```
CE3#traceroute 172.16.4.1 source loopback 1
Type escape sequence to abort.
Tracing the route to 172.16.4.1
VRF info: (vrf in name/id, vrf out name/id)
  1 192.168.103.30 0 msec 0 msec 0 msec
  2 192.168.31.1 [MPLS: Labels 16014/42 Exp 0] 0 msec 0 msec 0 msec
  3 192.168.1.2 [MPLS: Labels 16017/42 Exp 0] 4 msec 0 msec 0 msec
  4 192.168.104.40 [MPLS: Label 42 Exp 0] 0 msec 0 msec 0 msec
  5 192.168.104.41 0 msec 0 msec *
CE3#
CE4#trace 172.16.3.1 source loopback 1
Type escape sequence to abort.
Tracing the route to 172.16.3.1
VRF info: (vrf in name/id, vrf out name/id)
  1 192.168.104.40 0 msec 0 msec 0 msec
  2 192.168.42.2 [MPLS: Labels 16029/16026 Exp 0] 0 msec 0 msec 0 msec
  3 192.168.1.1 [MPLS: Labels 16004/16026 Exp 0] 4 msec 0 msec 0 msec
  4 192.168.31.30 [MPLS: Label 16026 Exp 0] 4 msec 0 msec 0 msec
  5 192.168.103.31 0 msec 0 msec *
CE4#
```

Lab 3-1: Establish Overlapping and Common Services Layer 3 VPNs

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

Activity Objective

In this lab activity, you will establish overlapping and common services Layer 3 VPNs to support customer needs.

Note Students from two pods will work in a team. The CE routers in both pods are running Cisco IOS Software. The first pod in the team will work on the PE router that is running Cisco IOS XR Software, and the second pod in the team will work on the PE router that is running Cisco IOS XE Software. Students in the same team should coordinate their lab activities.

You will work on different Cisco routers that are running Cisco IOS (c2900), Cisco IOS XE (asr1001), and Cisco IOS XR (asr9k) Software. After completing this activity, you will be able to meet these objectives:

- Enable overlapping Layer 3 VPNs
- Enable common services Layer 3 VPNs

Visual Objective

The figure illustrates what you will accomplish in this activity.

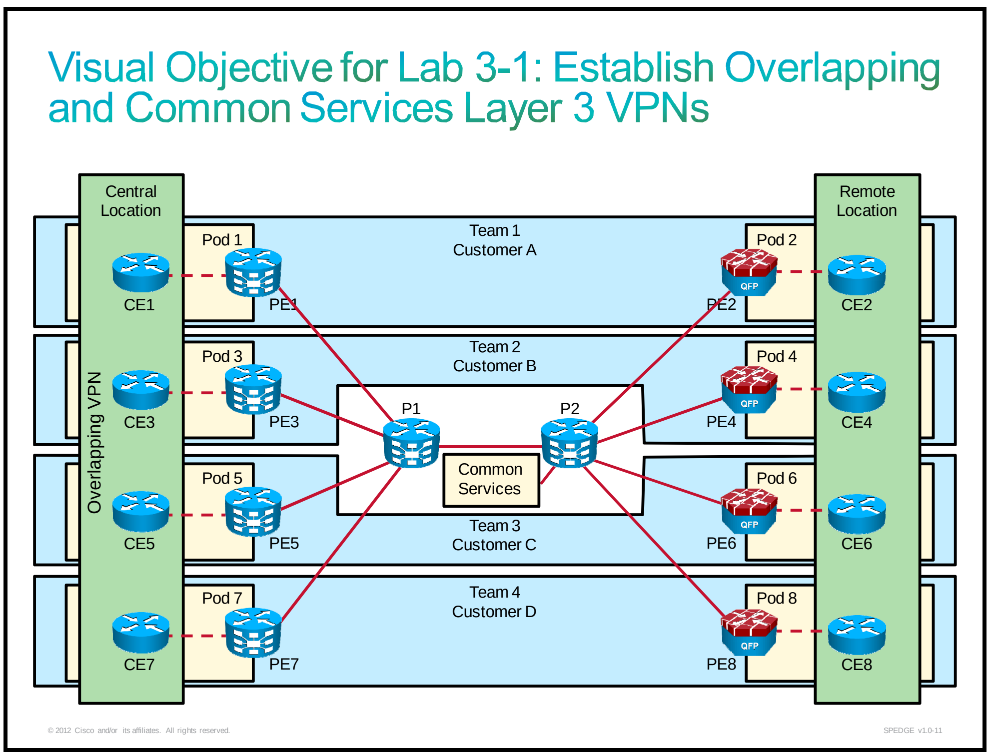

In the first task, you will establish overlapping VPNs that have the following connectivity goals:

- Simple VPN communication between the branch and central CE router (customer VPN) inside one team

- An overlapping VPN between central sites

There are also some common services that service providers offer. You will enable access to these services in Task 2.

VRF Assignments

This table gives the VRF RDs for VRFs that are used in this lab.

VRF Details

Team	Description	VRF RD
Team 1	Customer_1 Customer_1_C	1:210 1:211
Team 2	Customer_2 Customer_2_C	1:220 1:221
Team 3	Customer_3 Customer_3_C	1:230 1:231
Team 4	Customer_4 Customer_4_C	1:240 1:241
/	Common	1:1100

Required Resources

These resources and equipment are required to complete this activity:

- A PC with access to the Internet
- An SSH client that is installed on the PC

Command List

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

Cisco IOS and IOS XE Commands

Command	Description
configure terminal	Enters configuration mode
interface *interface*	Enters interface configuration mode
ip | ipv6 address *ip_address subnet_mask*	Sets a IPv4 or IPv6 address and the subnet mask on the interface
network *prefix wildcard_***mask area** *area*	Enables OSPF routing on the network and places the network into an OSPF area (in router configuration mode)
ping *dest_IP* **source** *source_IP*	Verifies connectivity between source IP and destination IP
router ospf *process-id*	Enables the OSPFv2 process on the router
router-id *IP-address*	Sets the OSPF router ID (in router configuration mode)
show ip interface brief	Displays the interface status and IPv4 addresses configured
show ip ospf	Displays information that is related to the OSPF routing protocol that is running on the router
show ip ospf interface	Displays OSPF interface information
show ip | ipv6 ospf database	Displays the contents of the OSPF database
show ip | ipv6 ospf neighbors	Displays OSPF neighbor information
show ip | ipv6 protocols	Displays IPv4 or IPv6 protocols that are running on the router.
show ip | ipv6 route	Displays the current routes in the routing table
show ip | ipv6 route vrf *vrf*	Displays the current routes in the VRF routing table
trace *dest_IP*	Traces the packet path through the network
router bgp *process-id*	Enables a BGP process on the router
neighbor *neighbor* **remote-as** *AS*	Assigns a BGP neighbor to an AS
neighbor *neighbor* **update-source** *Interface*	Sets the source interface for BGP updates
neighbor *neighbor* **activate**	Activates neighbors under address-family configuration mode
neighbor *neighbor* **send-community**	Sends communities with routing updates
neighbor *neighbor* **next-hop-self**	Updates the next-hop parameter for routing updates
redistribute *process process-id*	Redistributes route from other processes
vrf forwarding *vrf*	Assigns an interface to a specific VRF
ip vrf *vrf*	Configures a new VRF instance
rd *rd*	Sets a route distinguisher under VRF configuration
route-target export | import *rt*	Sets export and import route targets

Cisco IOS XR Commands

Command	Description
commit	Commits changes to the running configuration
configure terminal	Enters configuration mode
interface *interface*	Enters interface configuration mode (in global configuration mode)
interface *interface*	Defines the interfaces on which the OSPF protocol runs (in router configuration mode)
ipv4 | ipv6 address *ip_address/len*	Sets the IPv4 or IPv6 address for an interface and the subnet mask using the prefix length format
ipv6 enable	Enables IPv6 support on the interface
ping *dest_IP* **source** *source_IP*	Verifies connectivity between source IP and destination IP (IPv4 and IPv6)
router ospf | ospfv3 *process-ID*	Creates an OSPFv2 or OSPFv3 process
router-id *IP-address*	Sets an OSPF router ID (in router configuration mode)
network *prefix wildcard_***mask area** *area*	Enables OSPF routing on the network and places the network into an OSPF area (in router configuration mode)
show ipv4 interface brief	Displays the interface status and IPv4 addresses that are configured
show ospf	Displays information that is related to the OSPF routing protocol that is running on the router
show ospf interface	Displays OSPF interface information
show ospf | ospfv3 database	Displays the contents of the OSPF database
show ospf | ospfv3 neighbors	Displays OSPF neighbor information
show route	Displays the current routes in the routing table
show route vrf *vrf*	Displays the current routes in the VRF routing table
traceroute *IP-address*	Traces an IP address
router bgp *process-id*	Enables a BGP process on the router
neighbor *neighbor* **remote-as** *AS*	Assigns a BGP neighbor to an AS
neighbor *neighbor* **update-source** *Interface*	Sets the source interface for BGP updates
neighbor *neighbor* **address-family** *address-family*	Enters neighbor address family configuration mode
vrf *vrf*	Assigns an interface to a specific VRF (in interface configuration mode)
vrf *vrf*	Configures a new VRF instance (in global configuration mode)
rd *rd*	Set a route distinguisher under BGP VRF configuration
export | import route-target *rt*	Sets export and import route targets

Task 1: Enable Overlapping Layer 3 VPNs

In this task, you will enable connectivity between the central sites (Pod x) of each team. You will configure overlapping Layer 3 VPNs on the PEx routers.

Activity Procedure

Complete these steps to prepare the configuration for the routers in your pod:

Note	You will work with students from other pods and teams to finish this task. Students from different teams should coordinate their activities.

Step 1 You have to exchange BGP information with devices from other teams. Establish a BGP session with BGP route reflector P1 on address **10.0.1.1** and delete the neighbor relationship with the second PE router from your team. Use the BGP parameters that were used for the session with the second PE router. Check that the neighbor relationship is properly established.

Note	Most of the work in this task is performed by students in Pod x. Students from Pod y should help them.

Step 2 On the PEx router (only on that router), create new VRF called **Customer_z_C** (where z is your team number). Under the **ipv4 unicast address family** command, configure the router to export and import routes with RT **1:1000** and the RT of your team (**1:2z0**). Route target 1:1000 is used to exchange routes between the central sites.

Step 3 Reconfigure the OSPF process between the PEx and CEx routers. Replace Customer_z vrf with **Customer_z_C vrf**.

Step 4 Put the interface that is facing the CEx router into the Customer_z_C VRF. You will have to remove the IP address from the interface before you can change the VRF. Check connectivity with the CEx router. Check the OSPF neighbor relationship and routing table for VRF Customer_z_C.

Step 5 Reconfigure the BGP VRF section on router PEx. Replace Customer_z vrf with Customer_z_C vrf.

Activity Verification

Complete the verification of the lab activity:

- On the PE router in your pod, verify the BGP neighbors. On the PEx routers, verify the VRF Customer_z_C vrf routing table. On the PEy routers, verify the Customer_z VRF routing table. The PEx routers should also see routes from the central sites of other teams.

```
RP/0/RSP0/CPU0:PE3#sh route vrf Customer_2_C
<--- text omitted --->

Gateway of last resort is not set

B    172.16.1.1/32 [200/2] via 10.1.1.1 (nexthop in vrf default), 00:28:58
O    172.16.3.1/32 [110/2] via 192.168.103.31, 02:11:03,
GigabitEthernet0/0/0/0
B    172.16.4.1/32 [200/2] via 10.4.1.1 (nexthop in vrf default), 02:09:03
B    192.168.101.0/24 [200/0] via 10.1.1.1 (nexthop in vrf default), 00:28:58
C    192.168.103.0/24 is directly connected, 02:11:03, GigabitEthernet0/0/0/0
L    192.168.103.30/32 is directly connected, 02:11:03, GigabitEthernet0/0/0/0
```



```
B    192.168.104.0/24 [200/0] via 10.4.1.1 (nexthop in vrf default), 02:09:03
RP/0/RSP0/CPU0:PE3#
```

```
PE4#sh ip bgp summary
BGP router identifier 10.4.1.1, local AS number 64500
BGP table version is 7, main routing table version 7
2 network entries using 448 bytes of memory
2 path entries using 176 bytes of memory
2/2 BGP path/bestpath attribute entries using 416 bytes of memory
1 BGP community entries using 24 bytes of memory
1 BGP extended community entries using 60 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 1124 total bytes of memory
BGP activity 9/5 prefixes, 12/8 paths, scan interval 60 secs

Neighbor       V  AS    MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down State/PfxRcd
10.0.1.1       4 64500  1396    1497        7      0    0  22:34:44       2
PE4#
```

```
PE4#sh ip route vrf Customer_2

Routing Table: Customer_2
<--- text omitted --->

Gateway of last resort is not set

      172.16.0.0/32 is subnetted, 2 subnets
B        172.16.3.1 [200/2] via 10.3.1.1, 02:14:26
O        172.16.4.1 [110/2] via 192.168.104.41, 04:41:47, GigabitEthernet0/0/0
B     192.168.103.0/24 [200/0] via 10.3.1.1, 02:14:26
      192.168.104.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.104.0/24 is directly connected, GigabitEthernet0/0/0
L        192.168.104.40/32 is directly connected, GigabitEthernet0/0/0
PE4#
```

- On the CE router in your pod, verify the OSPF neighbors and routing table. There should be one OSPF neighbor, and you should get a default route by the OSPF process. Your output should be similar to the following:

```
CE1#sh ip route
<--- text omitted --->

Gateway of last resort is 192.168.101.10 to network 0.0.0.0

O*E2  0.0.0.0/0 [110/1] via 192.168.101.10, 00:35:06, GigabitEthernet0/1
      10.0.0.0/32 is subnetted, 1 subnets
C        10.1.10.1 is directly connected, Loopback0
      172.16.0.0/16 is variably subnetted, 4 subnets, 2 masks
C        172.16.1.0/24 is directly connected, Loopback1
L        172.16.1.1/32 is directly connected, Loopback1
      192.168.101.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.101.0/24 is directly connected, GigabitEthernet0/1
L        192.168.101.11/32 is directly connected, GigabitEthernet0/1
CE1#
```

```
CE1#sh ip ospf neighbor

Neighbor ID     Pri  State      Dead Time   Address         Interface
192.168.101.10    1  FULL/BDR   00:00:39    192.168.101.10  GigabitEthernet0/0
CE2#
```

- Verify connectivity between the central and branch site inside each team. Then verify connectivity to the central sites of other teams.

```
CE3#ping 172.16.4.1

Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 172.16.4.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
CE3#
CE3#ping 172.16.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
CE3#
CE4#ping 172.16.3.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.4.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
CE4#
CE4#ping 172.16.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.1.1, timeout is 2 seconds:
.....
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
CE4#
```

Task 2: Enable Common Services Layer 3 VPNs

In this task, you will establish connectivity to a common services device that is configured on the P2 router.

Activity Procedure

Complete this step:

Step 1 Reconfigure your customer VRF instance to exchange routes with the central services router. Import routes with RT **1:1102**. Export routes with RT **1:1101**.

Activity Verification

Complete the verification of the lab activity:

- Verify the VRF routing table on your PE router. You should see a route to network 172.16.100.100/32.

```
RP/0/RSP0/CPU0:PE3#sh route vrf Customer_2_C
<--- text omitted --->

Gateway of last resort is not set

B    172.16.1.1/32 [200/2] via 10.1.1.1 (nexthop in vrf default), 00:24:21
O    172.16.3.1/32 [110/2] via 192.168.103.31, 00:24:27,
GigabitEthernet0/0/0/0
B    172.16.4.1/32 [200/2] via 10.4.1.1 (nexthop in vrf default), 00:24:21
B    172.16.100.100/32 [200/0] via 10.0.2.1 (nexthop in vrf default), 00:11:10
B    192.168.101.0/24 [200/0] via 10.1.1.1 (nexthop in vrf default), 00:24:21
C    192.168.103.0/24 is directly connected, 00:24:28, GigabitEthernet0/0/0/0
L    192.168.103.30/32 is directly connected, 00:24:28, GigabitEthernet0/0/0/0
B    192.168.104.0/24 [200/0] via 10.4.1.1 (nexthop in vrf default), 00:24:21
RP/0/RSP0/CPU0:PE3#

PE4#sh ip route vrf Customer_2
<--- text omitted --->

Gateway of last resort is not set
```



```
      172.16.0.0/32 is subnetted, 3 subnets
B        172.16.3.1 [200/2] via 10.3.1.1, 00:36:50
O        172.16.4.1 [110/2] via 192.168.104.41, 22:17:27, GigabitEthernet0/0/0
B        172.16.100.100 [200/0] via 10.0.2.1, 00:02:39
B     192.168.103.0/24 [200/0] via 10.3.1.1, 00:36:50
      192.168.104.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.104.0/24 is directly connected, GigabitEthernet0/0/0
L        192.168.104.40/32 is directly connected, GigabitEthernet0/0/0
PE4#
```

- Verify connectivity to the central services server at IP address 172.16.100.100. Pinging should be successful.

```
CE4#ping 172.16.100.100

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.100.100, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

Lab 3-2: Establish Internet Connectivity with an MPLS Layer 3 VPN

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

Activity Objective

In this lab activity, you will establish Internet connectivity for a customer that also has an MPLS VPN tunnel established.

Note	Students from two pods will work in a team. The CE routers in both pods are running Cisco IOS Software. The first pod in the team will work on the PE router that is running Cisco IOS XR Software, and the second pod in the team will work on the PE router that is running Cisco IOS XE Software. Students in the same team should coordinate their lab activities.

You will work on different Cisco routers that are running Cisco IOS (c2900), Cisco IOS XE (asr1001), and Cisco IOS XR (asr9k) Software. After completing this activity, you will be able to meet these objectives:

- Establish PE-CE connectivity using subinterfaces
- Establish Internet connectivity
- Establish central site Internet connectivity
- Establish central site Internet connectivity across a separate MPLS VPN

Visual Objective

The figure illustrates what you will accomplish in this activity.

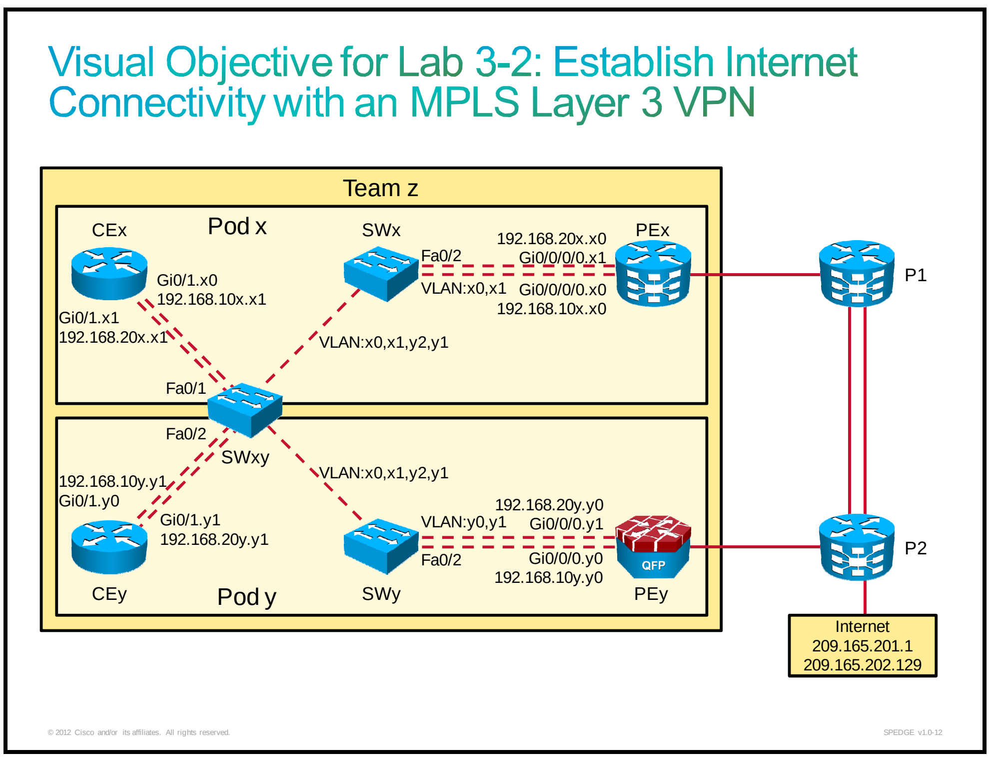

You will clear part of the configuration from the previous lab and configure two subinterfaces between the CE and PE routers. Using BGP, you will establish Internet connectivity for your site. In the next task, you will establish centralized Internet connectivity for remote sites through the central site. In the last task, you will use a separate MPLS VPN tunnel for Internet access.

The Internet is simulated on the P2 router. Two IP addresses are accessible for testing purposes (209.165.201.1 and 209.165.202.129). Routes are advertised in the global routing table and in the Internet VRF.

VRF Assignments

This table gives the VRF RDs for the VRFs that are used in this lab.

VRF Details

Team	Description	VRF RD
Team 1	Customer_1	1:210
Team 2	Customer_2	1:220
Team 3	Customer_3	1:230
Team 4	Customer_4	1:240
All	Internet	1:2000

Required Resources

These resources and equipment are required to complete this activity:

- A PC with access to the Internet
- An SSH client that is installed on the PC

Command List

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

Cisco IOS and IOS XE Commands

Command	Description
configure terminal	Enters configuration mode
interface *interface*	Enters interface configuration mode
interface *interface.subinterface*	Enters subinterface configuration mode
ip | ipv6 address *ip_address subnet_mask*	Sets a IPv4 or IPv6 address and the subnet mask on the interface
network *prefix wildcard_***mask area** *area*	Enables OSPF routing on the network and places the network into an OSPF area (in router configuration mode)
encapsulation dot1q *vlan*	Assigns a subinterface to the proper VLAN
ping *dest_IP* **source** *source_IP*	Verifies connectivity between source IP and destination IP
router ospf *process-id*	Enables an OSPFv2 process on the router
router-id *IP-address*	Sets an OSPF router ID (in router configuration mode)
show ip interface brief	Displays the interface status and IPv4 addresses configured
show ip ospf	Displays information that is related to the OSPF routing protocol that is running on the router
show ip ospf interface	Displays OSPF interface information
show ip | ipv6 ospf database	Displays the contents of the OSPF database
show ip | ipv6 ospf neighbors	Displays OSPF neighbor information
show ip | ipv6 protocols	Displays the IPv4 or IPv6 protocols that are running on the router.
show ip | ipv6 route	Displays the current routes in the routing table
show ip | ipv6 route vrf *vrf*	Displays the current routes in the VRF routing table
trace *dest_IP*	Traces a packet path through the network
router bgp *process-id*	Enables the BGP process on the router
neighbor *neighbor* **remote-as** *AS*	Assigns a BGP neighbor to an AS
neighbor *neighbor* **update-source** *Interface*	Sets the source interface for BGP updates
neighbor *neighbor* **activate**	Activates neighbors under address-family configuration mode
neighbor *neighbor* **send-community**	Sends communities with routing updates
neighbor *neighbor* **next-hop-self**	Updates the next-hop parameter for routing updates
redistribute *process process-id*	Redistributes route from other processes
vrf forwarding *vrf*	Assigns an interface to a specific VRF
ip vrf *vrf*	Configures a new VRF instance
rd *rd*	Sets a route distinguisher under VRF configuration

Command	Description
route-target export | import *rt*	Sets export and import route targets
switchport trunk allowed vlan vlan	Allows setting of VLANs across a trunk link
switchport mode trunk	Configures a switch port as a trunk port

Cisco IOS XR Commands

Command	Description
commit	Commits changes to the running configuration
configure terminal	Enters configuration mode
interface *interface*	Enters interface configuration mode (in global configuration mode)
interface *interface*	Defines the interfaces on which OSPF runs (in router configuration mode)
interface *interface.subinterface*	Enters subinterface configuration mode
ipv4 | ipv6 address *ip_address/len*	Sets the IPv4 or IPv6 address for an interface and the subnet mask using the prefix length format
ipv6 enable	Enables IPv6 support on the interface
encapsulation dot1q vlan	Assigns a subinterface in the proper VLAN
ping *dest_IP* **source** *source_IP*	Verifies connectivity between source IP and destination IP (IPv4 and IPv6)
router ospf | ospfv3 *process-ID*	Creates an OSPFv2 or OSPFv3 process
router-id *IP-address*	Sets an OSPF router ID (in router configuration mode)
network *prefix wildcard_***mask area** *area*	Enables OSPF routing on the network and places the network into an OSPF area (in router configuration mode)
show ipv4 interface brief	Displays the interface status and IPv4 addresses that are configured
show ospf	Displays information that is related to OSPF running on the router
show ospf interface	Displays OSPF interface information
show ospf | ospfv3 database	Displays the contents of the OSPF database
show ospf | ospfv3 neighbors	Displays OSPF neighbor information
show route	Displays the current routes in the routing table
show route vrf *vrf*	Displays the current routes in the VRF routing table
traceroute *IP-address*	Traces IP addresses
router bgp *process-id*	Enables the BGP process on the router
neighbor *neighbor* **remote-as** *AS*	Assigns a BGP neighbor to an AS
neighbor *neighbor* **update-source** *Interface*	Sets the source interface for BGP updates
neighbor *neighbor* **address-family** *address-family*	Enters neighbor address family configuration mode

Command	Description
vrf *vrf*	Assigns interface to specific VRF (in interface configuration mode)
vrf *vrf*	Configures a new VRF instance (in global configuration mode)
rd *rd*	Sets a route distinguisher under BGP VRF configuration
export | import route-target *rt*	Sets export and import route targets

Task 1: Restore a Simple Customer VPN Configuration

In this task, you will restore a simple customer VPN configuration for a site in your pod.

Activity Procedure

Complete these steps to prepare the configuration for the routers in your pod.

Note	You will work with students from other pods and teams to finish this task. Students from different teams should coordinate their activities.

Step 1 Reconfigure the CE-PE link on the PEx router. Set the link back to the **Customer_z** VRF, where z is your team number. VRF Customer_z should have the proper export and import targets configured (defined in the "VRF Details" table).

Step 2 Test connectivity between the sites. Connectivity between the central site and the remote site should work.

Activity Verification

Complete the verification of the lab activity:

- On the PE router in your pod, verify the BGP neighbors. On the PE routers, verify the Customer_z VRF routing table.

```
PE4#sh ip bgp summary
BGP router identifier 10.4.1.1, local AS number 64500
BGP table version is 7, main routing table version 7
2 network entries using 448 bytes of memory
2 path entries using 176 bytes of memory
2/2 BGP path/bestpath attribute entries using 416 bytes of memory
1 BGP community entries using 24 bytes of memory
1 BGP extended community entries using 60 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 1124 total bytes of memory
BGP activity 9/5 prefixes, 12/8 paths, scan interval 60 secs

Neighbor        V  AS    MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down State/PfxRcd
10.0.1.1        4 64500  1396    1497        7      0    0  22:34:44       2
PE4#

PE4#sh ip route vrf Customer_2

Routing Table: Customer_2
<--- text omitted --->

Gateway of last resort is not set

      172.16.0.0/32 is subnetted, 2 subnets
B        172.16.3.1 [200/2] via 10.3.1.1, 02:14:26
O        172.16.4.1 [110/2] via 192.168.104.41, 04:41:47, GigabitEthernet0/0/0
B     192.168.103.0/24 [200/0] via 10.3.1.1, 02:14:26
      192.168.104.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.104.0/24 is directly connected, GigabitEthernet0/0/0
L        192.168.104.40/32 is directly connected, GigabitEthernet0/0/0
PE4#
```

- On the CE router in your pod, verify the OSPF neighbors and routing table. There should be one OSPF neighbor, and you should get routes from other sites by the OSPF process. Your output should be similar to the following:

```
CE3#sh ip route
<--- text omitted --->
```

```
      10.0.0.0/32 is subnetted, 1 subnets
C        10.3.10.1 is directly connected, Loopback0
      172.16.0.0/16 is variably subnetted, 3 subnets, 2 masks
C        172.16.3.0/24 is directly connected, Loopback1
L        172.16.3.1/32 is directly connected, Loopback1
O E2     172.16.4.1/32
           [110/2] via 192.168.103.30, 01:13:13, GigabitEthernet0/1
      192.168.103.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.103.0/24 is directly connected, GigabitEthernet0/1
L        192.168.103.31/32 is directly connected, GigabitEthernet0/1
O E2  192.168.104.0/24
           [110/1] via 192.168.103.30, 01:13:13, GigabitEthernet0/1
CE3#

CE3#sh ip ospf neighbor
Neighbor ID     Pri   State    Dead Time   Address         Interface
10.3.1.1          1   FULL/DR  00:00:33    192.168.103.30  GigabitEthernet0/1
CE3#
```

- Verify connectivity between the central site and branch site inside each team.

```
CE3#ping 172.16.4.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.4.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
CE3#
CE4#ping 172.16.3.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.4.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
CE4#
```

Task 2: Establish CE-PE Connectivity for Internet Access

In this task, you will establish connectivity to a common services device that is configured on the P2 router.

Activity Procedure

Complete these steps:

Step 1 Reconfigure the PE-CE connecting interface on the PE and CE routers in your pod. Remove the IP parameters from the interface configuration.

Step 2 Create two subinterfaces, one for MPLS VPN access and one for Internet access. Use the parameters shown in the table for the subinterface configuration.

Router	Subinterface	VLAN	VRF	IP Address
PEx	Gi0/0/0/0.x0 Gi0/0/0/0.x1	x0 x1	Customer_z /	192.168.10x.x0 192.168.20x.x0
PEy	Gi0/0/0.y0 Gi0/0/0.y1	y0 y1	Customer_z /	192.168.10y.y0 192.168.20y.y0
CEx	Gi0/1.x0 Gi0/1.x1	x0 x1	/	192.168.10x.x1 192.168.20x.x1
CEy	Gi0/1.y0 Gi0/1.y1	y0 y1	/	192.168.10y.y1 192.168.20y.y1

Step 3 Configure VLAN **x1** and VLAN **y1** on all metro switches in your team. Allow VLANs x1 and y1 on all trunk links. Work together with the students from the other pod when configuring switch SWxy.

Step 4 The switch ports on the metro switches are now configured as access ports. Configure the ports that are facing your PE and CE routers as trunk ports. VLANs x0 and x1 should be allowed on the trunk links that are facing the Pod x routers. VLANs y0 and y1 should be allowed on the trunk links that are facing the Pod y routers. Work together with students from other pod when configuring switch SWxy.

Step 5 Reconfigure the OSPF process on PEx. Add interface **GigabitEthernet0/0/0/0.30** in the Area 0 section of the OSPF process.

Step 6 Verify connectivity between the PE and CE routers over both subinterfaces. Verify that OSPF is established between the PE and CE routers.

Activity Verification

Complete the verification of the lab activity:

- Verify connectivity between the PE and CE routers over both subinterfaces using ping.

```
CE4#ping 192.168.104.40
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.104.40, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
CE4#ping 192.168.204.40
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.104.40, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
CE3#sh ip ospf neighbor
Neighbor ID  Pri   State     Dead Time   Address         Interface
10.3.1.1       1   FULL/BDR  00:00:38    192.168.103.30  GigabitEthernet0/1.30
CE3#
```

Task 3: Establish Internet Connectivity

In this task, you will establish Internet connectivity for your pod site. You will use BGP for route distribution. Routes will be advertised in a global routing table.

Activity Procedure

Complete these steps:

Step 1 Reconfigure the BGP routing process on the PE router. Enable **ipv4 unicast address-family** globally for the BGP routing process. Enable **ipv4 unicast address-family** under **neighbor 10.0.1.1** and configure the **next-hop-self** feature for this neighbor.

Step 2 On the PE router, configure a new BGP neighbor to establish an EBGP session with the CE router. On the PEx router, set **192.168.20x.x1** as the neighbor IP address, **6450x** as the AS number, and interface **GigabitEthernet0/0/0/0.x1** as the update-source interface. On the PEy router, set **192.168.20y.y1** as the neighbor IP address, **6450y** as the AS number, and interface **GigabitEthernet0/0/0/0.y1** as the update-source interface. Enable **ipv4 unicast address-family** and configure the **next-hop-self** feature for this neighbor.

Step 3 On the PE router, advertise the default route. Block all other specific routes to be sent to this neighbor.

Step 4 Configure BGP process **6450x** on Pod x (or **6450y** in Pod y) on the CE router and establish a BGP neighbor relationship with the PE router. On the CEx router, set **192.168.20x.x0** as the neighbor IP address and **64500** as the AS number. On the CEy router, set **192.168.20y.y0** as the neighbor IP address and **64500** as the AS number.

Step 5 Advertise networks **192.168.20x.0/24** and **172.16.x.0/24** (or **192.168.20y.0/24** and **172.16.y.0/24** in Pod y).

Note Local network 172.16.x.0 has to be advertised in this example, because you do not use NAT to translate addresses in a public address space. In real situations, you usually use NAT.

Step 6 Verify the IPv4 routing table on the PE router. List all BGP routes.

Step 7 Verify Internet connectivity by pinging IP addresses **209.165.201.1** and **209.165.202.129**. Pinging should be successful.

Activity Verification

Complete the verification of the lab activity:

- Verify the IPv4 routing table on the PE router. List all BGP routes. You should see all Internet routes and routes from BGP-enabled sites.

```
RP/0/RSP0/CPU0:PE3#sh ip route bgp
Tue Nov 22 13:03:11.781 UTC

B    172.16.3.0/24 [20/0] via 192.168.203.31, 00:20:37
B    172.16.4.0/24 [200/0] via 10.4.1.1, 00:21:44
B    192.168.104.0/24 [200/0] via 10.4.1.1, 00:21:44
B    192.168.204.0/24 [200/0] via 10.4.1.1, 00:20:24
B    209.165.201.0/27 [200/0] via 10.0.2.1, 00:23:46
B    209.165.202.128/27 [200/0] via 10.0.2.1, 00:23:46
RP/0/RSP0/CPU0:PE3#
PE4#sh ip route bgp
<--- text omitted --->

      172.16.0.0/24 is subnetted, 2 subnets
B        172.16.3.0 [200/0] via 10.3.1.1, 00:22:38
B        172.16.4.0 [20/0] via 192.168.204.41, 00:23:45
B     192.168.104.0/24 [20/0] via 192.168.204.41, 00:23:45
B     192.168.203.0/24 [200/0] via 10.3.1.1, 00:22:08
      209.165.201.0/27 is subnetted, 1 subnets
B        209.165.201.0 [200/0] via 10.0.2.1, 02:05:21
      209.165.202.0/27 is subnetted, 1 subnets
B        209.165.202.128 [200/0] via 10.0.2.1, 02:05:21
PE4#
```

- Verify Internet connectivity using ping.

```
CE4#ping 209.165.201.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 209.165.201.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
CE3#ping 209.165.202.129

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 209.165.202.129, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```


Task 4: Establish Central Site Connectivity for Internet Access

In this task, you will establish central site Internet connectivity in your team. The Pod x router is the central router that provides Internet connectivity. The Pod y router is the branch router that uses the central site Internet connection to access the Internet.

Activity Procedure

Complete these steps:

Step 1 On the CEy router, shut down the **GigabitEthernet0/1.y1** subinterface that is used for Internet access.

Step 2 On the CEx router, reconfigure the OSPF process to advertise the default route to the other neighbors.

Step 3 On the CEx router, advertise the local network of remote sites (**172.16.y.0/24**) to the BGP process so that return traffic from the Internet router can be routable. You usually use NAT on a firewall to translate local addresses to public IP addresses, but in this exercise, local networks will be advertised to BGP.

Step 4 On the PEx router, advertise the default route in the BGP process under the Customer_z VRF configuration. Use the **network** command under the address-family area.

Step 5 Reconfigure the OSPF process on the PEy router that runs in the Customer_z VRF. Enable default route origination.

Step 6 Verify the routing tables on all routers in your pod. Verify Internet connectivity using the **ping** command.

Activity Verification

Complete the verification of the lab activity:

- Verify the routing tables on all routers. The default route should be in the routing table.

```
CE3#sh ip route
<--- text omitted --->

B*     0.0.0.0/0 [20/0] via 192.168.203.30, 00:54:07
       10.0.0.0/32 is subnetted, 1 subnets
C         10.3.10.1 is directly connected, Loopback0
       172.16.0.0/16 is variably subnetted, 3 subnets, 2 masks
C         172.16.3.0/24 is directly connected, Loopback1
L         172.16.3.1/32 is directly connected, Loopback1
O E2      172.16.4.0/24
            [110/2] via 192.168.103.30, 00:19:22, GigabitEthernet0/1.30
       192.168.103.0/24 is variably subnetted, 2 subnets, 2 masks
C         192.168.103.0/24 is directly connected, GigabitEthernet0/1.30
L         192.168.103.31/32 is directly connected, GigabitEthernet0/1.30
O E2   192.168.104.0/24
            [110/1] via 192.168.103.30, 00:55:54, GigabitEthernet0/1.30
       192.168.203.0/24 is variably subnetted, 2 subnets, 2 masks
C         192.168.203.0/24 is directly connected, GigabitEthernet0/1.31
L         192.168.203.31/32 is directly connected, GigabitEthernet0/1.31
CE3#
RP/0/RSP0/CPU0:PE3#sh route vrf Customer_2
<--- text omitted --->

O*E2 0.0.0.0/0 [110/1] via 192.168.103.31, 00:27:02, GigabitEthernet0/0/0/0.30
O    172.16.3.0/24 [110/2] via 192.168.103.31, 00:19:52,
GigabitEthernet0/0/0/0.30
B    172.16.4.0/24 [200/2] via 10.4.1.1 (nexthop in vrf default), 00:20:21
C    192.168.103.0/24 is directly connected, 01:09:12,
GigabitEthernet0/0/0/0.30
```

```
L    192.168.103.30/32 is directly connected, 01:09:12,
GigabitEthernet0/0/0/0.30
B    192.168.104.0/24 [200/0] via 10.4.1.1 (nexthop in vrf default), 00:56:53
RP/0/RSP0/CPU0:PE3#
PE4#sh ip route vrf Customer_2
<--- text omitted --->

B*     0.0.0.0/0 [200/1] via 10.3.1.1, 00:15:33
       172.16.0.0/24 is subnetted, 2 subnets
B         172.16.3.0 [200/2] via 10.3.1.1, 00:20:39
O         172.16.4.0
            [110/2] via 192.168.104.41, 00:21:08, GigabitEthernet0/0/0.40
B      192.168.103.0/24 [200/0] via 10.3.1.1, 00:57:40
       192.168.104.0/24 is variably subnetted, 2 subnets, 2 masks
C         192.168.104.0/24 is directly connected, GigabitEthernet0/0/0.40
L         192.168.104.40/32 is directly connected, GigabitEthernet0/0/0.40
PE4#
CE4#sh ip route
<--- text omitted --->

O*E2  0.0.0.0/0 [110/1] via 192.168.104.40, 00:16:07, GigabitEthernet0/1.40
       10.0.0.0/32 is subnetted, 1 subnets
C         10.4.10.1 is directly connected, Loopback0
       172.16.0.0/16 is variably subnetted, 3 subnets, 2 masks
O E2      172.16.3.0/24
            [110/2] via 192.168.104.40, 00:21:13, GigabitEthernet0/1.40
C         172.16.4.0/24 is directly connected, Loopback1
L         172.16.4.1/32 is directly connected, Loopback1
O E2  192.168.103.0/24
            [110/1] via 192.168.104.40, 00:58:14, GigabitEthernet0/1.40
       192.168.104.0/24 is variably subnetted, 2 subnets, 2 masks
C         192.168.104.0/24 is directly connected, GigabitEthernet0/1.40
L         192.168.104.41/32 is directly connected, GigabitEthernet0/1.40
CE4#
```

- Verify Internet access using the **ping** command. Loopback1 should be the source interface for ping traffic.

```
CE4#ping 209.165.201.1 source loopback1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 209.165.201.1, timeout is 2 seconds:
Packet sent with a source address of 172.16.4.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
CE4#
CE4#ping 209.165.202.129 source loopback1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 209.165.202.129, timeout is 2 seconds:
Packet sent with a source address of 172.16.4.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
CE4#
```

Task 5: Establish Central Site Connectivity for Internet Access Across a Separate MPLS VPN

In this task, you will establish central site Internet connectivity across a separate MPLS VPN in your team. VPN is preconfigured on the P routers.

Activity Procedure

Complete these steps:

Step 1 Create a new VRF instance called **Internet** on the PEx router. Under **ipv4 unicast address-family**, import and export routes with RT **1:2000**.

Step 2 Under the BGP process, configure a new VRF called **Internet** with RD **1:2000**. Under the Internet VRF, enable **ipv4 unicast address-family**.

Step 3 Reconfigure the interface that is facing the CEx router that is used for Internet access. Assign the interface to the Internet VRF.

Step 4 Delete the neighbor that is facing the CEx router and configure a new neighbor with the same parameters under the VRF Internet area.

Step 5 Verify the routing tables on the PEx and CEx routers in your pod. Verify Internet connectivity using the **ping** command.

Activity Verification

Complete the verification of the lab activity:

- Verify the routing tables on all routers.

```
CE3#sh ip route
<--- text omitted --->

B*     0.0.0.0/0 [20/0] via 192.168.203.30, 00:08:42
       10.0.0.0/32 is subnetted, 1 subnets
C         10.3.10.1 is directly connected, Loopback0
       172.16.0.0/16 is variably subnetted, 3 subnets, 2 masks
C         172.16.3.0/24 is directly connected, Loopback1
L         172.16.3.1/32 is directly connected, Loopback1
O E2      172.16.4.0/24
             [110/2] via 192.168.103.30, 00:56:54, GigabitEthernet0/1.30
       192.168.103.0/24 is variably subnetted, 2 subnets, 2 masks
C         192.168.103.0/24 is directly connected, GigabitEthernet0/1.30
L         192.168.103.31/32 is directly connected, GigabitEthernet0/1.30
O E2   192.168.104.0/24
             [110/1] via 192.168.103.30, 01:33:26, GigabitEthernet0/1.30
       192.168.203.0/24 is variably subnetted, 2 subnets, 2 masks
C         192.168.203.0/24 is directly connected, GigabitEthernet0/1.31
L         192.168.203.31/32 is directly connected, GigabitEthernet0/1.31
CE3#
RP/0/RSP0/CPU0:PE3#sh route vrf Internet
<--- text omitted --->

B    172.16.3.0/24 [20/0] via 192.168.203.31, 00:10:01
B    172.16.4.0/24 [20/2] via 192.168.203.31, 00:10:01
C    192.168.203.0/24 is directly connected, 00:32:49,
GigabitEthernet0/0/0/0.31
L    192.168.203.30/32 is directly connected, 00:32:49,
GigabitEthernet0/0/0/0.31
B    209.165.201.0/27 [200/0] via 10.0.2.1 (nexthop in vrf default), 00:32:52
B    209.165.202.128/27 [200/0] via 10.0.2.1 (nexthop in vrf default),
00:32:52
RP/0/RSP0/CPU0:PE3#
```

- Verify Internet access using the **ping** command.

```
CE4#ping 209.165.201.1 source loopback1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 209.165.201.1, timeout is 2 seconds:
Packet sent with a source address of 172.16.4.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
CE4#ping 209.165.202.129 source loopback1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 209.165.202.129, timeout is 2 seconds:
Packet sent with a source address of 172.16.4.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
CE4#
```

Lab 3-3: Implement CSC

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

Activity Objective

In this lab, you will implement the CSC feature using MP-BGP labels and route distribution.

Note	Students from two pods will work in a team. The CE routers in both pods are running Cisco IOS Software. The first pod in the team will work on the PE router that is running Cisco IOS XR Software, and the second pod in the team will work on the PE router that is running Cisco IOS XE Software. Students in the same team should coordinate their lab activities.

You will work on different Cisco routers that are running Cisco IOS (c2900), Cisco IOS XE (asr1001), and Cisco IOS XR (asr9k) Software. After completing this activity, you will be able to meet these objectives:

- Restore simple connectivity between the PE and CE routers
- Configure the CSC feature using MP-BGP

Visual Objective

The figure illustrates what you will accomplish in this activity.

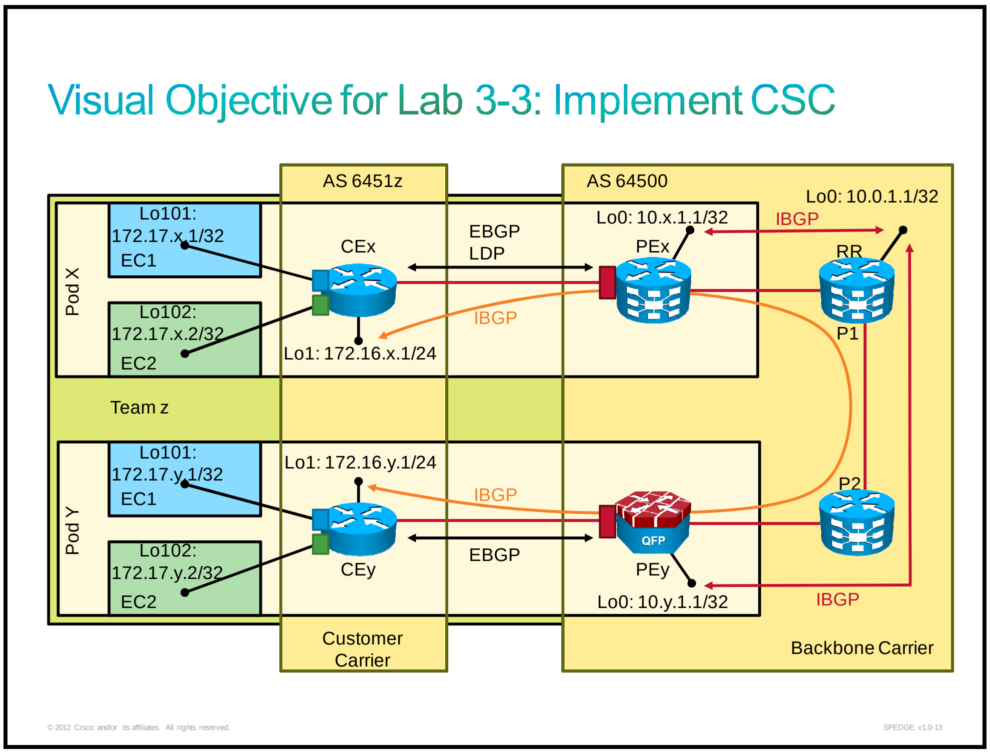

You will clear part of the configuration from the previous lab and restore a simple customer VPN configuration. Your new customer is a small service provider (customer carrier) that will use MPLS VPN service for interconnecting its point of presence (POP) sites.

The customer carrier has two customers, EC1 and EC2, that will be simulated using the Loopback101 and Loopback102 interfaces. Your goal is to establish MPLS VPN connectivity for customers EC1 and EC2 using the backbone carrier as the transport provider for the customer carrier.

VRF Assignments

This table gives the VRF RDs for the VRFs that are used in this lab.

Team	Name	VRF RD
Team 1	Customer_1	1:210
Team 2	Customer_2	1:220
Team 3	Customer_3	1:230
Team 4	Customer_4	1:240
All	EC1	1:321
All	EC2	1:322

Required Resources

These resources and equipment are required to complete this activity:

- A PC with access to the Internet
- An SSH client that is installed on the PC

Command List

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

Cisco IOS and IOS XE Commands

Command	Description
configure terminal	Enters configuration mode
interface *interface*	Enters interface configuration mode
interface *interface.subinterface*	Enters subinterface configuration mode
ip | ipv6 address *ip_address subnet_mask*	Sets a IPv4 or IPv6 address and the subnet mask on the interface
network *prefix wildcard_***mask area** *area*	Enables OSPF routing on the network and places a network into an OSPF area (in router configuration mode)
encapsulation dot1q vlan	Assigns a subinterface in the proper VLAN
ping *dest_IP* **source** *source_IP*	Verifies connectivity between source IP and destination IP
router ospf *process-id*	Enables an OSPFv2 process on the router.
router-id *IP-address*	Sets an OSPF router ID (in router configuration mode)
show ip interface brief	Displays the interface status and IPv4 addresses configured
show ip ospf	Displays information that is related to the OSPF routing protocol that is running on the router.
show ip ospf interface	Displays OSPF interface information
show ip | ipv6 ospf database	Displays the contents of the OSPF database
show ip | ipv6 ospf neighbors	Displays OSPF neighbor information
show ip | ipv6 protocols	Displays IPv4 or IPv6 protocols that are running on the router
show ip | ipv6 route	Displays the current routes in the routing table
show ip | ipv6 route vrf *vrf*	Displays the current routes in the VRF routing table
trace *dest_IP*	Traces a packet path through the network
router bgp *process-id*	Enables a BGP process on the router
neighbor *neighbor* **remote-as** *AS*	Assigns a BGP neighbor to an AS
neighbor *neighbor* **update-source** *Interface*	Sets the source interface for BGP updates
neighbor *neighbor* **activate**	Activates neighbors under address-family configuration mode
neighbor *neighbor* **send-community**	Sends communities with routing updates
neighbor *neighbor* **next-hop-self**	Updates the next-hop parameter for routing updates
redistribute *process process-id*	Redistributes routes from another process
vrf forwarding *vrf*	Assigns an interface to a specific VRF
ip vrf *vrf*	Configures a new VRF instance
rd *rd*	Sets a route distinguisher under VRF configuration

Command	Description
route-target export | import *rt*	Sets export and import route targets
switchport trunk allowed vlan vlan	Allows setting of VLANs across a trunk link
switchport mode trunk	Configures a switch port as a trunk port
address-family ipv4 labeled-unicast	Sends MPLS labels with BGP routes to a neighboring BGP router
allocate-label all	Allocates MPLS labels for specific IPv4 unicast or IPv6 unicast or VRF IPv4 unicast routes so that the BGP router can send labels with BGP routes to a neighboring router configured for labeled or VRF IPv6 unicast sessions (in the appropriate configuration mode). To restore the system to its default condition, use the **no** form of this command.

Cisco IOS XR Commands

Command	Description
commit	Commits changes to the running configuration.
configure terminal	Enters configuration mode
interface *interface*	Enters interface configuration mode (in global configuration mode)
interface *interface*	Defines the interfaces on which the OSPF protocol runs (in router configuration mode)
interface *interface.subinterface*	Enters subinterface configuration mode
ipv4 | ipv6 address *ip_address/len*	Sets the IPv4 or IPv6 address for an interface and the subnet mask using the prefix length format
ipv6 enable	Enables IPv6 support on the interface
encapsulation dot1q vlan	Assigns a subinterface in the proper VLAN
ping *dest_IP* **source** *source_IP*	Verifies connectivity between source IP and destination IP (IPv4 and IPv6)
router ospf | ospfv3 *process-ID*	Creates an OSPFv2 or OSPFv3 process
router-id *IP-address*	Sets an OSPF router ID (in router configuration mode)
network *prefix wildcard_***mask area** *area*	Enables OSPF routing on the network and places the network into an OSPF area (in router configuration mode)
show ipv4 interface brief	Displays the interface status and IPv4 addresses configured
show ospf	Displays information that is related to the OSPF routing protocol that is running on the router
show ospf interface	Displays OSPF interface information
show ospf | ospfv3 database	Displays the contents of the OSPF database
show ospf | ospfv3 neighbors	Displays OSPF neighbor information
show route	Displays the current routes in the routing table
show route vrf *vrf*	Displays the current routes in the VRF routing table

Command	Description
sh ip vrf interfaces	Displays the VRF interface mapping
traceroute *IP-address*	Traces an IP address
router bgp *process-id*	Enables a BGP process on the router
neighbor *neighbor* **remote-as** *AS*	Assigns a BGP neighbor to an AS
neighbor *neighbor* **send-label**	Sends MPLS labels with BGP routes to a neighboring BGP router
neighbor *neighbor* **update-source** *Interface*	Sets a source interface for BGP updates
neighbor *neighbor* **allowas-in**	Allows an AS path with the PE AS number
neighbor *neighbor* **address-family** *address-family*	Enters neighbor address family configuration mode
vrf *vrf*	Assigns an interface to specific VRF (in interface configuration mode)
vrf *vrf*	Configures a new VRF instance (in global configuration mode)
rd *rd*	Sets a route distinguisher under BGP VRF configuration
export | import route-target *rt*	Sets export and import route targets

Task 1: Restore Simple Connectivity Between the PE and CE Routers

In this task, you will restore a simple customer VPN configuration for a site in your pod.

Activity Procedure

Complete these steps to prepare the configuration for the routers in your pod:

Note You will work with students from other pods and teams to finish this task. Students from different teams should coordinate their activities.

Step 1 On all metro switches in your team, reconfigure all ports that are facing the PE and CE routers. Reconfigure the access ports and assign the proper VLAN to each port. Use the parameters in the table for port configuration.

Switch	Port	VLAN	Port Type
SWx	FastEthernet0/2	x0	nni
SWy	FastEthernet0/2	y0	nni
SWxy	FastEthernet0/1 FastEthernet0/2	x0 y0	nni nni

Step 2 Reconfigure the configuration of the interfaces between the PE and CE routers. Remove the subinterfaces and reconfigure the interfaces using the parameters that are shown in the table.

Router	Subinterface	VRF	IP Address
PEx	Gi0/0/0/0	Customer_z	192.168.10x.x0/24
PEy	Gi0/0/0	Customer_z	192.168.10y.y0/24
CEx	Gi0/1	/	192.168.10x.x1/24
CEy	Gi0/1	/	192.168.10y.y1/24

Step 3 Test connectivity between the PE and CE routers using the **ping** command.

Activity Verification

Complete the verification of the lab activity:

- Test connectivity between the PE and CE routers using the **ping** command.

```
CE3#ping 192.168.103.30
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.103.30, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

PE3#ping vrf Customer_2 192.168.103.31
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.103.30, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

Task 2: Simulate Customer Sites

In this task, you will configure end-customer VRF instances and loopback interfaces. Loopback interfaces will simulate the end-customer network.

Activity Procedure

Complete these steps:

Step 1 Configure two VRF instances on the CE routers. Use the parameters that are shown in the table for the VRF configuration.

Customer	VRF	Description	Import and Export Target	RD
Customer 1	EC1	End-customer 1 VRF	1:321	1:321
Customer 2	EC2	End-customer 2 VRF	1:322	1:322

Step 2 Create two loopback interfaces on the CE router that will be used for end-customer simulation. Use the parameters that are shown in the table for the subinterface configuration.

Router	Subinterface	VRF	IP Address
CEx	Loopback101 Loopback102	EC1 EC2	172.17.x.1/32 172.17.x.2/32
CEy	Loopback101 Loopback102	EC1 EC2	172.17.y.1/32 172.17.y.2/32

Activity Verification

Complete the verification of the lab activity:

- Verify the VRF EC1 and EC2 routing tables. Only directly connected routes should be present.

```
CE3#sh ip route vrf EC1
<--- text omitted --->

      172.17.0.0/32 is subnetted, 1 subnets
C        172.17.3.1 is directly connected, Loopback101
CE3#
CE3#sh ip route vrf EC2
<--- text omitted --->

      172.17.0.0/32 is subnetted, 1 subnets
C        172.17.3.2 is directly connected, Loopback102
CE3#
```

- Verify VRF assignments using the **show ip vrf interfaces** command.

```
CE3#sh ip vrf interfaces
Interface              IP-Address      VRF
Protocol
Lo101                  172.17.3.1      EC1                              up
Lo102                  172.17.3.2      EC2                              up
CE3#
```

Task 3: Configure Routing Between the PE and CE Routers

In this task, you will configure routing between the PE and CE routers using the BGP routing protocol. You will also use BGP for label distribution.

Activity Procedure

Complete these steps:

Step 1 Remove all BGP configuration on the CE router. Configure the BGP process with AS number **6451z**, where z is your team number. Establish an EBGP neighbor relationship between the PE and CE routers. Use the physical interface IP address for the source and destination of the BGP session. Use the **next-hop-self** command to change the next-hop parameter for routing updates.

Step 2 Change the subnet mask of the Loopback1 IP address on the CE routers to **/32**.

Step 3 On the CE router, advertise the network of the Loopback1 interface. This interface will be used to establish an IBGP session between the CE (customer carrier) routers. Reconfigure the PE routers to send labels with routing updates to both BGP neighbors. Use the **allocate-label** command to allocate labels for advertised prefixes.

Step 4 On Cisco IOS XR routers, an additional /32 static route to the CE router physical interface IP address (192.168.10x.x1) is required.

Task 4: Establish a BGP Session Between Customer Carrier Routers

In this task, you will establish an IBGP session between the customer carrier routers. You will establish MPLS VPN tunnels to carry traffic for end customers EC1 and EC2.

Activity Procedure

Complete these steps:

Step 1 Establish an IBGP session between the Loopback1 interfaces on the CE routers in your team. Use the Loopback1 interface IP addresses for the BGP source and destination addresses. Newly configured neighbors should be activated only under the VPNv4 address family.

Step 2 Configure the PE routers to override the AS number in the AS path for routing updates.

Step 3 Redistribute connected routes from VRF EC1 and VRF EC2.

Step 4 Verify connectivity between customer loopback addresses.

Activity Verification

Complete the verification of the lab activity:

- Verify the VRF EC1 and EC2 routing tables. Directly connected and BGP routes should be in the routing tables.

```
CE3#sh ip route vrf EC1
<--- text omitted --->
      172.17.0.0/32 is subnetted, 2 subnets
C        172.17.3.1 is directly connected, Loopback101
B        172.17.4.1 [200/0] via 172.16.4.1, 00:29:40
CE3#
```

- Verify connectivity between customer sites.

```
CE3#ping vrf EC1 172.17.4.1 so 172.17.3.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.17.4.1, timeout is 2 seconds:
Packet sent with a source address of 172.17.3.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
CE3#
```


Lab 4-1: Implement Layer 2 VPN (VPWS and VPLS)

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

Activity Objective

In this lab activity, you will configure an EoMPLS Layer 2 VPN tunnel between the CE routers in your team. You will use a pseudowire to interconnect two CE routers. Then you will configure a bridge group on one PE router and establish connectivity between CE routers in different teams.

Note	Students from two pods will work in a team. The CE routers in both pods are running Cisco IOS Software. The first pod in the team will work on the PE router that is running Cisco IOS XR Software, and the second pod in the team will work on the PE router that is running Cisco IOS XE Software. Students in the same team should coordinate their lab activities.

You will work on different Cisco routers that are running Cisco IOS (c2900), Cisco IOS XE (asr1001), and Cisco IOS XR (asr9k) Software. After completing this activity, you will be able to meet these objectives:

- Configure and monitor an xconnect on Cisco IOS XR Software
- Configure and monitor an xconnect on Cisco IOS XE Software
- Configure and monitor a bridge group and bridge domain

Visual Objective

The figure illustrates what you will accomplish in this activity.

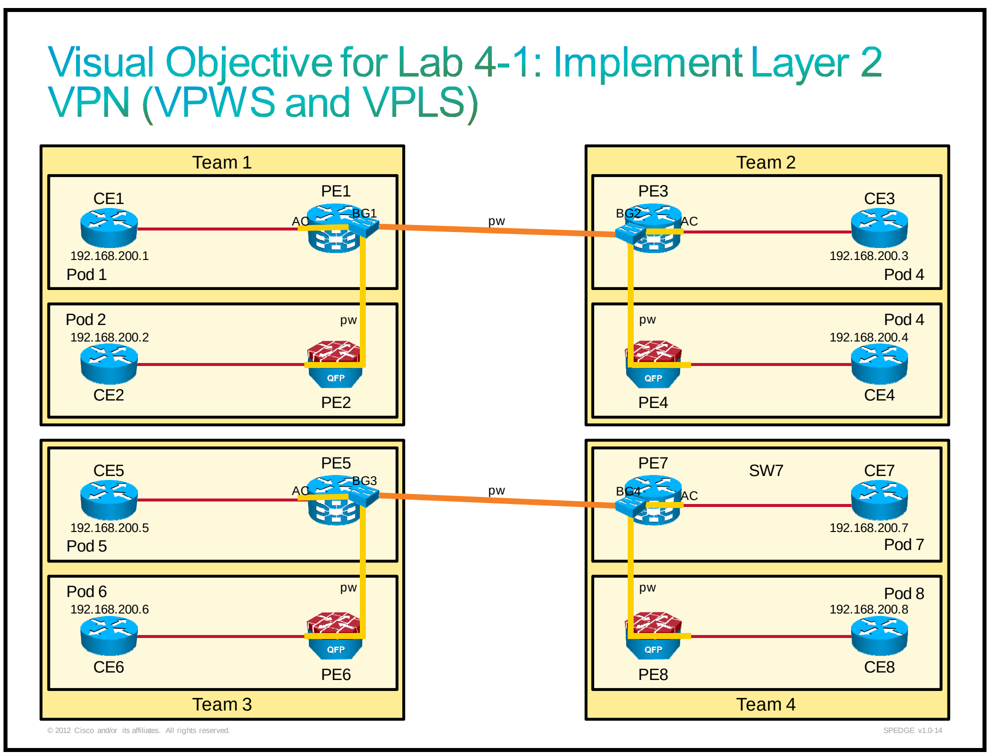

Command List

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

Cisco IOS XR Commands

Command	Description
l2vpn	Enters L2VPN configuration mode
pw-class	Defines a pseudowire class (in Layer 2 VPN configuration mode); assigns a class to a neighbor (in neighbor configuration mode)
encapsulation	Configures the encapsulation method (in pseudowire class configuration mode)
xconnect group	Defines a cross-connect group (in L2VPN configuration mode)
p2p	Enters p2p configuration submode to configure point-to-point cross-connects
interface	Attaches an interface to the VPWS (in p2p configuration submode)
neighbor	Configures a pseudowire for the VPWS (in p2p configuration submode)
interface *interface-id* **l2transport**	Enables an interface for the Layer 2 VPN transport service
rewrite ingress tag	Configures a VLAN tag rewrite operation for incoming and, optionally, outgoing frames
control-word	Enables the use of control words (in pseudowire class configuration mode)
transport-mode	Configures transport mode (in pseudowire class configuration mode)
bridge group	Defines a bridge group (in Layer 2 VPN configuration mode)
bridge domain	Defines a bridge domain (in Layer 2 VPN configuration mode)
vfi	Configures a virtual forwarding instance (in bridge domain configuration mode)
show l2vpn xconnect detail	Displays the xconnect parameters for the VPWS
sh l2vpn forwarding bridge-domain mac-address location	Displays learned MAC addresses
sh l2vpn bridge-domain [summary | detail]	Displays bridge domain parameters

Cisco IOS XE Commands

Command	Description
pseudowire-class	Defines a pseudowire class
encapsulation	Configures the encapsulation method (in pseudowire class configuration mode)
xconnect	Defines a VC (in interface configuration mode)
control-word	Enables the use of control words (in pseudowire class configuration mode)
show xconnect all detail	Displays the xconnect parameters for the VPWS

Task 1: Remove the CSC Configuration

In this task, you will remove the CSC configuration and prepare the lab for a new exercise.

Activity Procedure

Complete these steps:

Step 1 Remove the BGP configuration on the CE routers.

Step 2 Set the configuration of the interface between the CE and PE routers to the default values. Leave the speed and duplex configuration unchanged.

Activity Verification

Complete the verification of the lab activity:

- Verify that there is no BGP configuration on the CE routers. The PE-CE connecting interfaces should have the default configuration.

Task 2: Configure EoMPLS

In this task, you will configure a point-to-point EoMPLS tunnel for interconnecting two CE routers in your team.

Activity Procedure

Complete these steps:

Step 1 Configure the IP address of the PE-facing interface on the CE router. Use the IP address **192.168.200.x**, where x is your pod number. Use **/24** for the subnet mask.

Step 2 On the PE that is running Cisco IOS XR Software, configure the CE-facing interface for Layer 2 transport.

Step 3 On the PE running Cisco IOS XR Software, configure EoMPLS with these parameters:

- Suggested name for xconnect group: **podX-group** (where x is your pod number)
- Suggested name for point-to-point service: **podX-eompls** (where x is your pod number)
- Interface: the interface that is enabled for Layer 2 transport
- Neighbor: the Loopback0 address of the PE router in the neighbor pod
- Pseudowire ID: **12** (Team 1), **34** (Team 2), **56** (Team 3), **78** (Team 4)

- Pseudowire class: **podX-pw-class** (specify MPLS encapsulation)

Step 4 On the PE that is running Cisco IOS XE Software, configure a pseudowire class named **pw-class** that specifies MPLS encapsulation. On the CE-facing interface, configure an xconnect session. Use the Loopback0 interface of the neighbor PE router for the peer IP address, use the same pseudowire ID that you used in Step 3, and select **pw-class** for the pseudowire class.

Activity Verification

Complete the verification of the lab activity:

- Check the status of the Layer 2 VPN session. Examine its parameters. The xconnect status should be Up.

```
RP/0/RSP0/CPU0:PE3#sh l2vpn xconnect
Wed Feb  1 12:52:57.427 UTC
Legend: ST = State, UP = Up, DN = Down, AD = Admin Down, UR = Unresolved,
        SB = Standby, SR = Standby Ready

XConnect                   Segment 1                   Segment 2
Group      Name       ST   Description            ST   Description
ST
------------------------   -------------------------   -----------------------
xcon_grp   p2p        UP   Gi0/0/0/0              UP   10.4.1.1        1
UP
--------------------------------------------------------------------------------
RP/0/RSP0/CPU0:PE3#
RP/0/RSP0/CPU0:PE3#sh l2vpn xconnect detail
Wed Feb  1 12:52:48.095 UTC

Group xcon_grp, XC p2p, state is up; Interworking none
  AC: GigabitEthernet0/0/0/0, state is up
    Type Ethernet
    MTU 1500; XC ID 0x840001; interworking none
    Statistics:
      packets: received 1293018, sent 822336
      bytes: received 149458334, sent 475654155
  PW: neighbor 10.4.1.1, PW ID 1, state is up ( established )
    PW class pod3-pw-class, XC ID 0x840001
    Encapsulation MPLS, protocol LDP
    PW type Ethernet, control word enabled, interworking none
    PW backup disable delay 0 sec
    Sequencing not set

      MPLS         Local                          Remote
      ------------ ------------------------------ ---------------------------
      Label        16024                          41
      Group ID     0x4000080                      0x0
      Interface    GigabitEthernet0/0/0/0         unknown
      MTU          1500                           1500
      Control word enabled                        enabled
      PW type      Ethernet                       Ethernet
      VCCV CV type 0x2                            0x2
                   (LSP ping verification)        (LSP ping verification)
      VCCV CC type 0x7                            0x7
                   (control word)                 (control word)
                   (router alert label)           (router alert label)
                   (TTL expiry)                   (TTL expiry)
      ------------ ------------------------------ ---------------------------
    MIB cpwVcIndex: 4294705153
    Create time: 01/02/2012 12:30:10 (00:22:36 ago)
    Last time status changed: 01/02/2012 12:30:15 (00:22:31 ago)
    Statistics:
      packets: received 822336, sent 1293018
      bytes: received 475654155, sent 149458334
RP/0/RSP0/CPU0:PE3#
```

```
PE4#sh xconnect all detail
Legend:    XC ST=Xconnect State  S1=Segment1 State  S2=Segment2 State
  UP=Up       DN=Down            AD=Admin Down      IA=Inactive
  SB=Standby  HS=Hot Standby     RV=Recovering      NH=No Hardware

XC ST  Segment 1                         S1 Segment 2
S2
------+---------------------------------+--+--------------------------------
UP     ac   Gi0/0/0(Ethernet)            UP mpls 10.3.1.1:1
UP
            Interworking: none                   Local  VC label 41
                                                 Remote VC label 16024
                                                 pw-class: pw-class1
```

- From your CE router, ping the CE router in the partner pod. The ping should work if the pseudowire was established successfully.

```
CE3#ping 192.168.200.4
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.200.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
CE3#
```

Task 3: Configure VPLS

In this task, you will configure a bridge domain on the PE router that is running the Cisco IOS XR operating system. You will establish connectivity between Team 1 and Team 2 and between Team 3 and Team 4.

Activity Procedure

Complete these steps:

Step 1 Reconfigure the PE router running the Cisco IOS XR operating system. Remove the xconnect group and configure a new bridge group named **BG1**. Under BG1, configure a new bridge domain named **BD1**.

Step 2 Configure interface GigabitEthernet0/0/0/0 as a member of BD1.

Step 3 Under BD1, configure two pseudowires. Use the parameters from the table.

PE	Pseudowire Neighbor 1	Pseudowire ID 1	Pseudowire Neighbor 2	Pseudowire ID 2
PE1	10.2.1.1	12	10.3.1.1	13
PE3	10.4.1.1	34	10.1.1.1	13
PE5	10.6.1.1	56	10.7.1.1	57
PE7	10.8.1.1	78	10.5.1.1	57

Activity Verification

Complete the verification of the lab activity:

- Check the bridge-domain parameters.

```
RP/0/RSP0/CPU0:PE1#sh l2vpn bridge-domain summary
Thu Feb  2 23:55:45.524 PST
Number of groups: 1, bridge-domains: 1, Up: 1, Shutdown: 0
Default: 1, pbb-edge: 0, pbb-core: 0
Number of ACs: 1 Up: 1, Down: 0
```

```
Number of PWs: 2 Up: 2, Down: 0, Standby: 0
RP/0/RSP0/CPU0:PE1#

RP/0/RSP0/CPU0:PE1#sh l2vpn bridge-domain detail
Thu Feb  2 23:59:56.786 PST
Bridge group: BG1, bridge-domain: BD1, id: 0, state: up, ShgId: 0, MSTi: 0
  MAC learning: enabled
  MAC withdraw: enabled
    MAC withdraw for Access PW: enabled
  Flooding:
    Broadcast & Multicast: enabled
    Unknown unicast: enabled
  MAC aging time: 300 s, Type: inactivity
  MAC limit: 4000, Action: none, Notification: syslog
  MAC limit reached: no
  MAC port down flush: enabled
  MAC Secure: disabled, Logging: disabled
  Split Horizon Group: none
  Dynamic ARP Inspection: disabled, Logging: disabled
  IP Source Guard: disabled, Logging: disabled
  DHCPv4 snooping: disabled
  IGMP Snooping profile: none
  Bridge MTU: 1500
  MIB cvplsConfigIndex: 1
  Filter MAC addresses:
  Create time: 02/02/2012 01:41:27 (22:18:29 ago)
  No status change since creation
  ACs: 1 (1 up), VFIs: 0, PWs: 2 (2 up), PBBs: 0 (0 up)
  List of ACs:
    AC: GigabitEthernet0/0/0/0, state is up
      Type Ethernet
      MTU 1500; XC ID 0x840001; interworking none
      MAC learning: enabled
      Flooding:
        Broadcast & Multicast: enabled
        Unknown unicast: enabled
      MAC aging time: 300 s, Type: inactivity
      MAC limit: 4000, Action: none, Notification: syslog
      MAC limit reached: no
      MAC port down flush: enabled
      MAC Secure: disabled, Logging: disabled
      Split Horizon Group: none
      Dynamic ARP Inspection: disabled, Logging: disabled
      IP Source Guard: disabled, Logging: disabled
      DHCPv4 snooping: disabled
      IGMP Snooping profile: none
      Storm Control: disabled
      Static MAC addresses:
      Statistics:
        packets: received 6113833, sent 4835644
        bytes: received 981542339, sent 579406455
      Storm control drop counters:
        packets: broadcast 0, multicast 0, unknown unicast 0
        bytes: broadcast 0, multicast 0, unknown unicast 0
      Dynamic ARP inspection drop counters:
        packets: 0, bytes: 0
      IP source guard drop counters:
        packets: 0, bytes: 0
  List of Access PWs:
    PW: neighbor 10.2.1.1, PW ID 10, state is up ( established )
      PW class pod1-pw-class, XC ID 0xfffc0002
      Encapsulation MPLS, protocol LDP
      PW type Ethernet, control word disabled, interworking none
      PW backup disable delay 0 sec
      Sequencing not set

        MPLS         Local                          Remote
        ------------ ------------------------------ -------------------------
```



```
      Label        16026                          44
      Group ID     0x0                            0x0
      Interface    Access PW                      unknown
      MTU          1500                           1500
      Control word disabled                       disabled
      PW type      Ethernet                       Ethernet
      VCCV CV type 0x2                            0x2
                   (LSP ping verification)        (LSP ping verification)
      VCCV CC type 0x6                            0x6
                   (router alert label)           (router alert label)
                   (TTL expiry)                   (TTL expiry)
      ------------ ------------------------------ --------------------------
    MIB cpwVcIndex: 0
    Create time: 02/02/2012 23:55:09 (00:04:47 ago)
    Last time status changed: 02/02/2012 23:55:14 (00:04:42 ago)
    MAC withdraw message: send 0 receive 0
    Static MAC addresses:
    Statistics:
      packets: received 45, sent 153
      bytes: received 6864, sent 12085
    Storm control drop counters:
      packets: broadcast 0, multicast 0, unknown unicast 0
      bytes: broadcast 0, multicast 0, unknown unicast 0
    MAC learning: enabled
    Flooding:
      Broadcast & Multicast: enabled
      Unknown unicast: enabled
    MAC aging time: 300 s, Type: inactivity
    MAC limit: 4000, Action: none, Notification: syslog
    MAC limit reached: no
    MAC port down flush: enabled
    MAC Secure: disabled, Logging: disabled
    Split Horizon Group: none
    DHCPv4 snooping: disabled
    IGMP Snooping profile: none
    Storm Control: disabled
  PW: neighbor 10.3.1.1, PW ID 13, state is up ( established )
    PW class pod1-pw-class, XC ID 0xfffc0003
    Encapsulation MPLS, protocol LDP
    PW type Ethernet, control word disabled, interworking none
    PW backup disable delay 0 sec
    Sequencing not set

      MPLS         Local                          Remote
      ------------ ------------------------------ --------------------------
      Label        16020                          16027
      Group ID     0x0                            0x0
      Interface    Access PW                      vfi2
      MTU          1500                           1500
      Control word disabled                       disabled
      PW type      Ethernet                       Ethernet
      VCCV CV type 0x2                            0x2
                   (LSP ping verification)        (LSP ping verification)
      VCCV CC type 0x6                            0x6
                   (router alert label)           (router alert label)
                   (TTL expiry)                   (TTL expiry)
      ------------ ------------------------------ --------------------------
    MIB cpwVcIndex: 4294705155
    Create time: 02/02/2012 23:55:09 (00:04:47 ago)
    Last time status changed: 02/02/2012 23:55:22 (00:04:34 ago)
    MAC withdraw message: send 0 receive 0
    Static MAC addresses:
    Statistics:
      packets: received 9, sent 153
      bytes: received 2152, sent 14318
    Storm control drop counters:
      packets: broadcast 0, multicast 0, unknown unicast 0
      bytes: broadcast 0, multicast 0, unknown unicast 0
```

```
      MAC learning: enabled
      Flooding:
        Broadcast & Multicast: enabled
        Unknown unicast: enabled
      MAC aging time: 300 s, Type: inactivity
      MAC limit: 4000, Action: none, Notification: syslog
      MAC limit reached: no
      MAC port down flush: enabled
      MAC Secure: disabled, Logging: disabled
      Split Horizon Group: none
      DHCPv4 snooping: disabled
      IGMP Snooping profile: none
      Storm Control: disabled
  List of VFIs:
RP/0/RSP0/CPU0:PE1#
```

- Check connectivity between the CE routers. You should ping all the CE routers in both teams.

```
CE1#ping 192.168.200.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.200.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
CE1#
CE1#ping 192.168.200.3
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.200.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
CE1#
CE1#ping 192.168.200.4
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.200.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
CE1#
```

- Check the MAC address forwarding table.

```
PE1#sh l2vpn forwarding bridge-domain mac-address location 0/0/CPU0
Mac Address    Type    Learned from/Filtered  LC learned Resync Age Mapped to
-------------------------------------------------------------------------------
4055.3986.f969 dynamic (10.2.1.1, 10)         0/0/CPU0   0d 0h 0m 4s     N/A
5835.d9d6.0204 dynamic (10.3.1.1, 13)         0/0/CPU0   0d 0h 0m 0s     N/A
e8b7.482c.a211 dynamic (10.3.1.1, 13)         0/0/CPU0   0d 0h 0m 0s     N/A
e8ba.70b5.5e04 dynamic (10.2.1.1, 10)         0/0/CPU0   0d 0h 0m 0s     N/A
e8ba.70b5.6404 dynamic Gi0/0/0/0              0/0/CPU0   0d 0h 0m 0s     N/A
```

Task 4: Use a VFI

In this task, you will configure a VFI and move the pseudowire configuration under it. The VFI implements a split-horizon rule, so traffic between the two pseudowires within the VFI is dropped to prevent loops.

Activity Procedure

Complete these steps:

Step 1 Under bridge domain BD1, configure a VFI named **vfi1**.

Step 2 Move all pseudowire configuration under the VFI area.

Activity Verification

Because a full mesh of pseudowires is not established within Team 1 and 2 nor within Team 3 and 4 in this exercise to illustrate VFI split-horizon operations, after successful completion of the VFI configurations, you should see these results:

- CE1 pings to CE2 and to CE3 should be successful. Pings from CE1 to CE4 should fail, for these reasons:
  - — PE1 has pseudowires only to PE2 and PE3.
  - — PE1 does not have a pseudowire to PE4.
- CE2 pings to CE1 should be successful. Pings from CE2 to CE3 and to CE4 should fail, for these reasons:
  - — PE2 has a pseudowire only to PE1.
  - — PE2 does not have a pseudowire to PE3 or to PE4.

As with CE1 and CE2, you should see the same ping behavior on CE3 to CE8.

- Check the bridge-domain parameters.

```
RP/0/RSP0/CPU0:PE1#sh l2vpn bridge-domain summary
Thu Feb  2 23:39:25.417 PST
Number of groups: 1, bridge-domains: 1, Up: 1, Shutdown: 0
Default: 1, pbb-edge: 0, pbb-core: 0
Number of ACs: 1 Up: 1, Down: 0
Number of PWs: 2 Up: 2, Down: 0, Standby: 0
RP/0/RSP0/CPU0:PE1#

RP/0/RSP0/CPU0:PE1#sh l2vpn bridge-domain detail
Thu Feb  2 23:41:02.599 PST
Bridge group: BG1, bridge-domain: BD1, id: 0, state: up, ShgId: 0, MSTi: 0
  MAC learning: enabled
  MAC withdraw: enabled
    MAC withdraw for Access PW: enabled
  Flooding:
    Broadcast & Multicast: enabled
    Unknown unicast: enabled
  MAC aging time: 300 s, Type: inactivity
  MAC limit: 4000, Action: none, Notification: syslog
  MAC limit reached: no
  MAC port down flush: enabled
  MAC Secure: disabled, Logging: disabled
  Split Horizon Group: none
  Dynamic ARP Inspection: disabled, Logging: disabled
  IP Source Guard: disabled, Logging: disabled
  DHCPv4 snooping: disabled
  IGMP Snooping profile: none
  Bridge MTU: 1500
  MIB cvplsConfigIndex: 1
  Filter MAC addresses:
  Create time: 02/02/2012 01:41:27 (21:59:35 ago)
  No status change since creation
  ACs: 1 (1 up), VFIs: 1, PWs: 2 (2 up), PBBs: 0 (0 up)
  List of ACs:
    AC: GigabitEthernet0/0/0/0, state is up
      Type Ethernet
      MTU 1500; XC ID 0x840001; interworking none
      MAC learning: enabled
      Flooding:
        Broadcast & Multicast: enabled
        Unknown unicast: enabled
      MAC aging time: 300 s, Type: inactivity
      MAC limit: 4000, Action: none, Notification: syslog
      MAC limit reached: no
```

```
      MAC port down flush: enabled
      MAC Secure: disabled, Logging: disabled
      Split Horizon Group: none
      Dynamic ARP Inspection: disabled, Logging: disabled
      IP Source Guard: disabled, Logging: disabled
      DHCPv4 snooping: disabled
      IGMP Snooping profile: none
      Storm Control: disabled
      Static MAC addresses:
      Statistics:
        packets: received 6113115, sent 4835563
        bytes: received 981492354, sent 579386043
      Storm control drop counters:
        packets: broadcast 0, multicast 0, unknown unicast 0
        bytes: broadcast 0, multicast 0, unknown unicast 0
      Dynamic ARP inspection drop counters:
        packets: 0, bytes: 0
      IP source guard drop counters:
        packets: 0, bytes: 0
  List of Access PWs:
  List of VFIs:
    VFI vfi1
      PW: neighbor 10.2.1.1, PW ID 10, state is up ( established )
        PW class pod1-pw-class, XC ID 0xfffc0002
        Encapsulation MPLS, protocol LDP
        PW type Ethernet, control word disabled, interworking none
        PW backup disable delay 0 sec
        Sequencing not set

          MPLS         Local                          Remote
          ------------ ------------------------------ -----------------------
          Label        16026                          41
          Group ID     0x0                            0x0
          Interface    vfi1                           unknown
          MTU          1500                           1500
          Control word disabled                       disabled
          PW type      Ethernet                       Ethernet
          VCCV CV type 0x2                            0x2
                       (LSP ping verification)        (LSP ping verification)
          VCCV CC type 0x6                            0x6
                       (router alert label)           (router alert label)
                       (TTL expiry)                   (TTL expiry)
          ------------ ------------------------------ -----------------------
        MIB cpwVcIndex: 0
        Create time: 02/02/2012 01:58:20 (21:42:42 ago)
        Last time status changed: 02/02/2012 01:58:25 (21:42:37 ago)
        MAC withdraw message: send 1 receive 0
        Static MAC addresses:
        Statistics:
          packets: received 9316, sent 40304
          bytes: received 981541, sent 2838677
      DHCPv4 snooping: disabled
      IGMP Snooping profile: none
      PW: neighbor 10.3.1.1, PW ID 13, state is up ( established )
        PW class pod1-pw-class, XC ID 0xfffc0003
        Encapsulation MPLS, protocol LDP
        PW type Ethernet, control word disabled, interworking none
        PW backup disable delay 0 sec
        Sequencing not set

          MPLS         Local                          Remote
          ------------ ------------------------------ -----------------------
          Label        16020                          16027
          Group ID     0x0                            0x0
          Interface    vfi1                           vfi2
          MTU          1500                           1500
          Control word disabled                       disabled
          PW type      Ethernet                       Ethernet
```



```
    VCCV CV type 0x2                              0x2
                 (LSP ping verification)          (LSP ping verification)
    VCCV CC type 0x6                              0x6
                 (router alert label)             (router alert label)
                 (TTL expiry)                     (TTL expiry)
    ------------ ------------------------------ -------------------------
  MIB cpwVcIndex: 4294705155
  Create time: 02/02/2012 01:59:34 (21:41:27 ago)
  Last time status changed: 02/02/2012 02:01:44 (21:39:18 ago)
  MAC withdraw message: send 0 receive 0
  Static MAC addresses:
  Statistics:
    packets: received 1475, sent 40241
    bytes: received 507754, sent 2833022
DHCPv4 snooping: disabled
IGMP Snooping profile: none
VFI Statistics:
  drops: illegal VLAN 0, illegal length 0
RP/0/RSP0/CPU0:PE1#
```

- Check connectivity between the CE routers. Where the split-horizon rule is applied, pinging fails.

```
CE1#ping 192.168.200.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.200.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
CE1#
CE1#
CE1#ping 192.168.200.3
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.200.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
CE1#
CE1#ping 192.168.200.4
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.200.4, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
CE1#

CE2#ping 192.168.200.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.200.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
CE2#
CE2#
CE2#ping 192.168.200.3
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.200.3, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
CE2#
CE2#
CE2#ping 192.168.200.4
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.200.4, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
CE2#
```

- Check the MAC address forwarding table.

```
PE1#sh l2vpn forwarding bridge-domain mac-address location 0/0/CPU0
Mac Address    Type    Learned from/Filtered  LC learned Resync Age Mapped to
--------------------------------------------------------------------------------
5835.d9d6.0204 dynamic (10.3.1.1, 13)         0/0/CPU0   0d 0h 0m 0s    N/A
e8b7.482c.a181 dynamic Gi0/0/0/0              0/0/CPU0   0d 0h 0m 0s    N/A
e8ba.70b5.5e04 dynamic (10.2.1.1, 10)         0/0/CPU0   0d 0h 0m 0s    N/A
e8ba.70b5.6404 dynamic Gi0/0/0/0              0/0/CPU0   0d 0h 0m 0s    N/A
```


Answer Key

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

Lab 2-1 Answer Key: Implement MPLS Layer 3 VPN Backbones

When you complete this lab activity, the device configuration and device outputs will be similar to the results shown here, with differences that are specific to your pod.

Task 1: Configure VRF Tables Necessary to Support the Customer

Step 1 Enable LDP:

PEx (Cisco IOS XR):

```
mpls ldp
 interface gigabitEthernet 0/0/0/2
!
```

PEy (Cisco IOS XE):

```
mpls label protocol ldp
!
interface GigabitEthernet0/0/3
 mpls ip
!
```

Step 2 Configure a VRF instance:

PEx (Cisco IOS XR):

```
vrf Customer_1
 address-family ipv4 unicast
  import route-target 1:210
  export route-target 1:210
!
```

PEy (Cisco IOS XE):

```
ip vrf Customer_1
 rd 1:210
 route-target export 1:210
 route-target import 1:210
!
```

Step 3 Reconfigure the interface:

PEx (Cisco IOS XR):

```
interface gigabitEthernet0/0/0/0
 no ipv4 address
 vrf Customer_1
 ipv4 address 192.168.101.10 255.255.255.0
!
```

PEy (Cisco IOS XE):

```
interface gigabitEthernet0/0/0
 ip vrf forwarding Customer_1
 ip address 192.168.102.20 255.255.255.0
!
```

Task 2: Configure MP-BGP to Establish Routing Between the PE Routers

Step 1 Enable a BGP process:

PEx (Cisco IOS XR):

```
router bgp 64500
 neighbor 10.0.1.1
  remote-as 64500
```

```
  update-source Loopback0
!
```

PEy (Cisco IOS XE):

```
router bgp 64500
 neighbor 10.0.1.1 remote-as 64500
 neighbor 10.0.1.1 update-source Loopback0
 !
```

Step 2 Enable a BGP process:

PEx (Cisco IOS XR):

```
router bgp 64500
 address-family vpnv4 unicast
 !
 neighbor 10.0.1.1
   address-family vpnv4 unicast
   next-hop-self
!
```

PEy (Cisco IOS XE):

```
router bgp 64500
 address-family vpnv4
  neighbor 10.0.1.1 activate
  neighbor 10.0.1.1 send-community both
  neighbor 10.0.1.1 next-hop-self
 exit-address-family
!
```

Lab 2-2 Answer Key: Connect Customers to MPLS Layer 3 VPNs

When you complete this lab activity, the device configuration and device outputs will be similar to the results shown here, with differences that are specific to your pod.

Task 1: Configuring Static Routes Between the PE and CE Routers

Step 1 Configure a loopback interface:

CEx (Cisco IOS):

```
interface Loopback1
 ip address 172.16.3.1 255.255.255.0
!
```

CEy (Cisco IOS):

```
interface Loopback1
 ip address 172.16.4.1 255.255.255.0
!
```

Step 2 Configure a default route on the CE router:

CEx (Cisco IOS):

```
ip route 0.0.0.0 0.0.0.0 192.168.103.30
!
```

CEy (Cisco IOS):

```
ip route 0.0.0.0 0.0.0.0 192.168.104.40
!
```

Step 3 Configure a static route on the PE router:

PEx (Cisco IOS XR):

```
router static
 vrf Customer_2
  address-family ipv4 unicast
   172.16.3.0/24 192.168.103.31
  !
```



```
!
```

PEy (Cisco IOS XE):

```
ip route vrf Customer_2 172.16.4.0 255.255.255.0 192.168.104.41
!
```

Step 4 Configure route redistribution:

PEx (Cisco IOS XR):

```
router bgp 64500
 vrf Customer_2
  rd 1:220
  address-family ipv4 unicast
   redistribute connected
   redistribute static
!
```

PEy (Cisco IOS XE):

```
router bgp 64500
 address-family ipv4 vrf Customer_2
  redistribute connected
  redistribute static
 exit-address-family
!
```

Task 2: Configure RIP as the PE-CE Routing Protocol

Step 1 Remove the static routes:

CEx (Cisco IOS):

```
no ip route 0.0.0.0 0.0.0.0 192.168.103.30
!
```

CEy (Cisco IOS):

```
no ip route 0.0.0.0 0.0.0.0 192.168.104.40
!
```

PEx (Cisco IOS XR):

```
router static
 vrf Customer_2
  address-family ipv4 unicast
   no 172.16.3.0/24 192.168.103.31
!
```

PEy (Cisco IOS XE):

```
no ip route vrf Customer_2 172.16.4.0 255.255.255.0 192.168.104.41
!
```

Step 2 Configure RIP between the CE and PE routers:

CEx (Cisco IOS):

```
router rip
 version 2
 network 172.16.0.0
 network 192.168.103.0
 no auto-summary
!
```

CEy (Cisco IOS):

```
router rip
 version 2
 network 172.16.0.0
 network 192.168.104.0
 no auto-summary!
```

PEx (Cisco IOS XR):

```
router rip
 vrf Customer_2
  interface GigabitEthernet0/0/0/0
```

```
  !
!
```

PEy (Cisco IOS XE):

```
router rip
 version 2
 !
 address-family ipv4 vrf Customer_2
  network 192.168.104.0
  no auto-summary
  version 2
 exit-address-family
!
```

Step 3 Redistribute RIP routes:

PEx (Cisco IOS XR):

```
router bgp 64500
 vrf Customer_2
  address-family ipv4 unicast
   no redistribute static
   redistribute rip
  !
 !
!
```

PEy (Cisco IOS XE):

```
router bgp 64500
 address-family ipv4 vrf Customer_2
  no redistribute static
  redistribute rip
 exit-address-family
!
```

Step 4 Redistribute BGP routes:

PEx (Cisco IOS XR):

```
router rip
 vrf Customer_2
  redistribute bgp 64500
 !
!
```

PEy (Cisco IOS XE):

```
router rip
 address-family ipv4 vrf Customer_2
  redistribute bgp 64500 metric transparent
 exit-address-family
!
```

Task 3: Configure EIGRP as the PE-CE Routing Protocol

Step 1 Configure EIGRP:

CEx (Cisco IOS):

```
no router rip
!
router eigrp 1
 network 172.16.3.0 0.0.0.255
 network 192.168.103.0
!
```

CEy (Cisco IOS):

```
no router rip
!
router eigrp 1
```



```
 network 172.16.4.0 0.0.0.255
 network 192.168.104.0
!
```

PEx (Cisco IOS XR):

```
no router rip
!
router eigrp 1
 vrf Customer_2
  address-family ipv4
   autonomous-system 1
   interface GigabitEthernet0/0/0/0
   !
!
```

PEy (Cisco IOS XE):

```
no router rip
!
router eigrp 1
 !
 address-family ipv4 vrf Customer_2 autonomous-system 1
  network 192.168.104.0
 exit-address-family
!
```

Step 2 Redistribute EIGRP routes:

PEx (Cisco IOS XR):

```
router bgp 64500
 vrf Customer_2
   address-family ipv4 unicast
   no redistribute rip
   redistribute eigrp 1
  !
!
```

PEy (Cisco IOS XE):

```
router bgp 64500
 address-family ipv4 vrf Customer_2
  no redistribute rip
  redistribute eigrp 1
 exit-address-family
!
```

Step 3 Redistribute BGP routes:

PEx (Cisco IOS XR):

```
router eigrp 1
 vrf Customer_2
  address-family ipv4
   default-metric 10000 100 255 1 1500
   redistribute bgp 64500
   !
```

PEy (Cisco IOS XE):

```
router eigrp 1
 address-family ipv4 vrf Customer_2 autonomous-system 1
  redistribute bgp 64500 metric 10000 100 255 1 1500
 exit-address-family
!
```

Lab 2-3 Answer Key: Connect Advanced Customers to MPLS Layer 3 VPNs

When you complete this lab activity, the device configuration and device outputs will be similar to the results shown here, with differences that are specific to your pod.

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

Step 1 Remove the EIGRP configuration and configure a BGP process:

CEx (Cisco IOS):

```
no router eigrp 1
!
router bgp 64503
!
```

CEy (Cisco IOS):

```
no router eigrp 1
!
router bgp 64504
!
```

PEx (Cisco IOS XR):

```
no router eigrp 1
!
router bgp 64500
 vrf Customer_2
   address-family ipv4 unicast
   no redistribute eigrp 1
  !
```

PEy (Cisco IOS XE):

```
no router eigrp 1
!
router eigrp 1
 address-family ipv4 vrf Customer_2 autonomous-system 1
  no redistribute bgp 64500
exit-address-family
!
```

Step 2 Configure **route-policy**:

PEx (Cisco IOS):

```
route-policy pass
  pass
end-policy
!
```

Step 3 Configure a BGP neighbor relationship:

CEx (Cisco IOS):

```
router bgp 64503
 neighbor 192.168.103.30 remote-as 64500
 neighbor 192.168.103.30 update-source GigabitEthernet0/1
 neighbor 192.168.103.30 next-hop-self!
```

CEy (Cisco IOS):

```
router bgp 64504
 neighbor 192.168.104.40 remote-as 64500
 neighbor 192.168.104.40 update-source GigabitEthernet0/1
 neighbor 192.168.104.40 next-hop-self!
```

PEx (Cisco IOS XR):

```
router bgp 64500
 vrf Customer_2
```



```
  neighbor 192.168.103.31
   remote-as 64503
   update-source GigabitEthernet0/0/0/0
   address-family ipv4 unicast
    next-hop-self
   !
```

PEy (Cisco IOS XE):

```
router bgp 64500
 address-family ipv4 vrf Customer_2
  neighbor 192.168.104.41 remote-as 64504
  neighbor 192.168.104.41 update-source GigabitEthernet0/0/0
  neighbor 192.168.104.41 activate
  neighbor 192.168.104.41 next-hop-self
 exit-address-family
```

Step 4 Advertise routes:

CEx (Cisco IOS):

```
router bgp 64503
 network 172.16.3.0 mask 255.255.255.0
!
```

CEy (Cisco IOS):

```
router bgp 64504
 network 172.16.4.0 mask 255.255.255.0
!
```

PEx (Cisco IOS XR):

```
router bgp 64500
 vrf Customer_2
  neighbor 192.168.103.31
   address-family ipv4 unicast
    route-policy pass in
    route-policy pass out
   !
  !
```

Task 2: Configure OSPF as the PE-CE Routing Protocol

Step 1 Remove EBGP and configure OSPF:

CEx (Cisco IOS):

```
no router bgp 64503
!
router ospf 220
!
```

CEy (Cisco IOS):

```
no router bgp 64504
!
router ospf 220
!
```

PEx (Cisco IOS XR):

```
router bgp 64500
 vrf Customer_2
  neighbor 192.168.103.31
   shutdown
!
```

PEy (Cisco IOS XE):

```
router ospf 220 vrf Customer_2
!
```

Step 2 Advertise routes:

CEx (Cisco IOS):

```
interface Loopback1
 ip ospf network point-to-point
!
router ospf 220
 network 172.16.3.0 0.0.0.255 area 0
 network 192.168.103.0 0.0.0.255 area 0
!
```

CEy (Cisco IOS):

```
interface Loopback1
 ip ospf network point-to-point
!
router ospf 220
 network 172.16.4.0 0.0.0.255 area 0
 network 192.168.104.0 0.0.0.255 area 0
!
```

PEx (Cisco IOS XR):

```
router ospf 220
 vrf Customer_2
  address-family ipv4 unicast
  area 0
   interface GigabitEthernet0/0/0/0
!
```

PEy (Cisco IOS XE):

```
router ospf 220 vrf Customer_2
 network 192.168.104.0 0.0.0.255 area 0
```

Step 3 Redistribute OSPF routes into BGP:

PEx (Cisco IOS XR):

```
router bgp 64500
 vrf Customer_2
  address-family ipv4 unicast
   redistribute ospf 220
  !
!
```

PEy (Cisco IOS XE):

```
router bgp 64500
 address-family ipv4 vrf Customer_2
  redistribute ospf 220
!
```

Step 4 Redistribute BGP routes into OSPF:

PEx (Cisco IOS XR):

```
router bgp 64500
 vrf Customer_2
  address-family ipv4 unicast
   redistribute ospf 220
  !
!
```

PEy (Cisco IOS XE):

```
router ospf 220
 vrf Customer_2
  redistribute bgp 64500
!
```

Lab 3-1 Answer Key: Establish Overlapping and Common Services Layer 3 VPNs

When you complete this lab activity, the device configuration and device outputs will be similar to the results shown here, with differences that are specific to your pod.

Task 1: Enable Overlapping Layer 3 VPNs

Step 1 Reconfigure the BGP process and change the peer router:

PE4 (Cisco IOS):

```
router bgp 64500
 neighbor 10.0.1.1 remote-as 64500
 neighbor 10.0.1.1 update-source Loopback0
 address-family vpnv4
  neighbor 10.0.1.1 activate
  neighbor 10.0.1.1 send-community both
  neighbor 10.0.1.1 next-hop-self
 exit-address-family
```

PE3 (Cisco IOS XR):

```
router bgp 64500
 neighbor 10.0.1.1
  remote-as 64500
  update-source Loopback0
  address-family vpnv4 unicast
!
neighbor 10.4.1.1
  shutdown
!
!
```

Step 2 Reconfigure the VRF instance:

PE3 (Cisco IOS XR):

```
vrf Customer_2_C
 description VPN Customer 2 Central
 address-family ipv4 unicast
  import route-target
   1:220
   1:1000
  !
  export route-target
   1:220
   1:1000
  !
 !
!
```

Step 3 Reconfigure the OSPF process:

PE3 (Cisco IOS XR):

```
router ospf 102
 vrf Customer_2_C
  default-information originate always
  address-family ipv4 unicast
  area 0
   interface GigabitEthernet0/0/0/0
   !
  !
 !
!
```

Step 4 Put an interface into the new VRF:

PE3 (Cisco IOS XR):

```
interface GigabitEthernet0/0/0/0
 no ipv4 address
 !
 vrf Customer_2_C
 ipv4 address 192.168.103.30 255.255.255.0
 speed 100
!
```

Step 5 Enable IPv6 on the CE and PE routers:

PE3 (Cisco IOS XR):

```
router bgp 64500
 no vrf Customer_2
 !
 vrf Customer_2_C
  rd 1:202
  address-family ipv4 unicast
   redistribute ospf 102
  !
 !
!
```

Task 2: Enable Common Services Layer 3 VPNs

Step 1 Reconfigure the VRF instance:

PE3 (Cisco IOS XR):

```
vrf Customer_2_C
 address-family ipv4 unicast
  import route-target
   1:1102
  !
  export route-target
   1:1101
  !
 !
!
```

PE4 (Cisco IOS):

```
ip vrf Customer_2
 route-target export 1:1101
 route-target import 1:1102
!
```

Lab 3-2 Answer Key: Establish Internet Connectivity with an MPLS Layer 3 VPN

When you complete this lab activity, the device configuration and device outputs will be similar to the results shown here, with differences that are specific to your pod.

Task 1: Restore a Simple Customer VPN Configuration

Step 1 Reconfigure the CE-PE link:

PE3 (Cisco IOS XR):

```
interface GigabitEthernet0/0/0/0
 no ipv4 address 192.168.103.30 255.255.255.0
 no vrf Customer_2_C
 vrf Customer_2
 ipv4 address 192.168.103.30 255.255.255.0
!
vrf Customer_2
 description VPN Customer 2
 address-family ipv4 unicast
  import route-target
   1:220
```



```
  !
  export route-target
   1:220
  !
  maximum prefix 1000
 !
!
```

Task 2: Establish CE-PE Connectivity for Internet Access

Step 1 Reconfigure the VRF instance:

PE3 (Cisco IOS XR):

```
interface GigabitEthernet0/0/0/0
 no ipv4 address
 no vrf Customer_2
```

PE4 (Cisco IOS):

```
interface GigabitEthernet0/0/0
 no ip address
 no ip vrf forwarding Customer_2
```

CE3 (Cisco IOS):

```
interface GigabitEthernet0/1
 no ip address
```

CE4 (Cisco IOS):

```
interface GigabitEthernet0/1
 no ip address
```

Step 2 Create subinterfaces:

PE3 (Cisco IOS XR):

```
interface GigabitEthernet0/0/0/0.30
 vrf Customer_2
 ipv4 address 192.168.103.30 255.255.255.0
 encapsulation dot1q 30
!
interface GigabitEthernet0/0/0/0.31
 ipv4 address 192.168.203.30 255.255.255.0
 encapsulation dot1q 31
```

PE4 (Cisco IOS):

```
interface GigabitEthernet0/0/0.40
 encapsulation dot1Q 40
 ip vrf forwarding Customer_2
 ip address 192.168.104.40 255.255.255.0
!
interface GigabitEthernet0/0/0.41
 encapsulation dot1Q 41
 ip address 192.168.204.40 255.255.255.0
!
```

CE3 (Cisco IOS):

```
interface GigabitEthernet0/1.30
 encapsulation dot1Q 30
 ip address 192.168.103.31 255.255.255.0
!
interface GigabitEthernet0/1.31
 encapsulation dot1Q 31
 ip address 192.168.203.31 255.255.255.0
```

```
!
```

CE4 (Cisco IOS):

```
interface GigabitEthernet0/1.40
 encapsulation dot1Q 40
 ip address 192.168.104.41 255.255.255.0
!
interface GigabitEthernet0/1.41
 encapsulation dot1Q 41
 ip address 192.168.204.41 255.255.255.0
```

Step 3 Configure VLANs:

SW3 (Cisco IOS):

```
vlan 31
 name vlan31
!
vlan 41
 name vlan41
!
interface FastEthernet0/23
 switchport trunk allowed vlan add 31,41
!
```

SW4 (Cisco IOS):

```
vlan 31
 name vlan31
!
vlan 41
 name vlan41
!
interface FastEthernet0/23
 switchport trunk allowed vlan add 31,41
!
```

SW34 (Cisco IOS):

```
vlan 31
 name vlan31
!
vlan 41
 name vlan41
!
interface FastEthernet0/21
 switchport trunk allowed vlan add 31,41
!
interface FastEthernet0/23
 switchport trunk allowed vlan add 31,41
!
```

Step 4 Configure switch ports:

SW3 (Cisco IOS):

```
interface FastEthernet0/2
 switchport trunk allowed vlan 30,31
 switchport mode trunk
!
```

SW4 (Cisco IOS):

```
interface FastEthernet0/2
 switchport trunk allowed vlan 40,41
 switchport mode trunk
!
```

SW34 (Cisco IOS):

```
interface FastEthernet0/1
 switchport trunk allowed vlan 30,31
 switchport mode trunk
```



```
!
interface FastEthernet0/2
 switchport trunk allowed vlan 40,41
 switchport mode trunk
!
```

Step 5 Reconfigure OSPF on the PE routers:

PE3 (Cisco IOS XR):

```
router ospf 102
 vrf Customer_2
  area 0
   interface GigabitEthernet0/0/0/0.30
!
```

Task 3: Establish Internet Connectivity

Step 1 Reconfigure the BGP routing process on the PE routers:

PE3 (Cisco IOS XR):

```
router bgp 64500
 address-family ipv4 unicast
 !
neighbor 10.0.1.1
  address-family ipv4 unicast
   next-hop-self
```

PE4 (Cisco IOS):

```
router bgp 64500
address-family ipv4
  neighbor 10.0.1.1 activate
  neighbor 10.0.1.1 next-hop-self
```

Step 2 Configure a new BGP peer:

PE3 (Cisco IOS XR):

```
router bgp 64500
 neighbor 192.168.203.31
  remote-as 64503
  update-source GigabitEthernet0/0/0/0.31
  address-family ipv4 unicast
   next-hop-self
  !
```

PE4 (Cisco IOS):

```
router bgp 64500
 neighbor 192.168.204.41 remote-as 64504
 address-family ipv4
  neighbor 192.168.204.41 activate
  neighbor 192.168.204.41 next-hop-self
```

Step 3 Configure a new BGP peer on the PE routers:

PE3 (Cisco IOS XR):

```
router bgp 64500
 neighbor 192.168.203.31
  address-family ipv4 unicast
   route-policy pass in
   route-policy Only_Default out
   default-originate
  !
```

```
route-policy pass
  pass
end-policy
!
route-policy Only_Default
  if destination in (0.0.0.0/0) then
    pass
  endif
end-policy
!
```

PE4 (Cisco IOS):

```
router bgp 64500
 address-family ipv4
  neighbor 192.168.204.41 default-originate
  neighbor 192.168.204.41 prefix-list Only_Default out
!
ip prefix-list Only_Default seq 5 permit 0.0.0.0/0
!
```

Step 4 Configure a new BGP peer on the CE routers:

CE3 (Cisco IOS):

```
router bgp 64503
 neighbor 192.168.203.30 remote-as 64500
!
```

CE4 (Cisco IOS):

```
router bgp 64504
 neighbor 192.168.204.40 remote-as 64500
!
```

Step 5 Advertise networks on the CE routers:

CE3 (Cisco IOS):

```
router bgp 64503
 network 172.16.3.0 mask 255.255.255.0
 network 192.168.203.0
!
```

CE4 (Cisco IOS):

```
router bgp 64504
 network 172.16.4.0 mask 255.255.255.0
 network 192.168.204.0
!
```

Task 4: Establish Central Site Connectivity for Internet Access

Step 1 Shut down the subinterface:

CE4 (Cisco IOS):

```
interface GigabitEthernet0/1.41
 shutdown
!
```

Step 2 Reconfigure OSPF on the CEx router:

CE3 (Cisco IOS):

```
router ospf 1
```



```
 default-information originate
```

Step 3 Reconfigure BGP on the CEx router:

CE3 (Cisco IOS):

```
router bgp 64503
 network 172.16.4.0 mask 255.255.255.0
!
```

Step 4 Reconfigure BGP on the PEx router:

PE3 (Cisco IOS XR):

```
router bgp 64500
 vrf Customer_2
  address-family ipv4 unicast
   network 0.0.0.0/0
  !
```

Step 5 Reconfigure OSPF on the PEy router:

PE4 (Cisco IOS):

```
router ospf 102 vrf Customer_2
 default-information originate
!
```

Task 5: Establish Central Site Connectivity for Internet Access Across a Separate MPLS VPN

Step 1 Create a new VRF instance:

PE3 (Cisco IOS XR):

```
vrf Internet
 description Internet Simulation
 address-family ipv4 unicast
  import route-target
   1:2000
  !
  export route-target
   1:2000
  !
```

Step 2 Reconfigure the BGP process:

PE3 (Cisco IOS XR):

```
router bgp 64500
vrf Internet
  rd 1:2000
  address-family ipv4 unicast
!
```

Step 3 Reconfigure the subinterface:

PE3 (Cisco IOS XR):

```
interface GigabitEthernet0/0/0/0.31
 no ipv4 address
 vrf Internet
 ipv4 address 192.168.203.30 255.255.255.0
!
```

Step 4 Reconfigure the BGP neighbor:

PE3 (Cisco IOS XR):

```
router bgp 64500
 neighbor 192.168.203.31
 shutdown
!
router bgp 64500
 vrf Internet
  neighbor 192.168.203.31
   remote-as 64503
   update-source GigabitEthernet0/0/0/0.31
   address-family ipv4 unicast
    route-policy pass in
    route-policy Only_Default out
    default-originate
    next-hop-self
!
```

Lab 3-3 Answer Key: Implement CSC

When you complete this lab activity, the device configuration and device outputs will be similar to the results shown here, with differences that are specific to your pod.

Task 1: Restore Simple Connectivity Between PE and CE Routers

Step 1 Reconfigure switch ports on the metro switches:

SW1:

```
interface FastEthernet0/2
 port-type nni
 switchport access vlan 10
 switchport mode access
```

SW2:

```
interface FastEthernet0/2
 port-type nni
 switchport access vlan 20
 switchport mode access
```

SW12:

```
interface FastEthernet0/1
 port-type nni
 switchport access vlan 10
 switchport mode access
!
interface FastEthernet0/2
 port-type nni
 switchport access vlan 20
 switchport mode access
```

Step 2 Reconfigure interfaces between the PE and CE routers:

PE1 (Cisco IOS XR):

```
no interface GigabitEthernet0/0/0/0.10
no interface GigabitEthernet0/0/0/0.11
!
interface GigabitEthernet0/0/0/0
 no ipv4 address
 vrf Customer_1
 ipv4 address 192.168.101.10 255.255.255.0
!
```

CE1 (Cisco IOS):

```
no interface GigabitEthernet0/1.10
no interface GigabitEthernet0/1.11
```



```
!
interface GigabitEthernet0/1
 ip address 192.168.101.11 255.255.255.0
```

PE2 (Cisco IOS XE):

```
no interface GigabitEthernet0/0/0.10
no interface GigabitEthernet0/0/0.11
!
interface GigabitEthernet0/0/0
 ip vrf forwarding Customer_1
 ip address 192.168.102.20 255.255.255.0
```

CE2 (Cisco IOS):

```
no interface GigabitEthernet0/1.10
no interface GigabitEthernet0/1.11
!
interface GigabitEthernet0/1
 ip address 192.168.102.21 255.255.255.0
```

Task 2: Simulate Customer Sites

Step 1 Reconfigure the VRF instance:

CE1 (Cisco IOS):

```
ip vrf EC1
 rd 1:321
 route-target export 1:321
 route-target import 1:321
!
ip vrf EC2
 rd 1:322
 route-target export 1:322
 route-target import 1:322
!
```

CE2 (Cisco IOS):

```
ip vrf EC1
 rd 1:321
 route-target export 1:321
 route-target import 1:321
!
ip vrf EC2
 rd 1:322
 route-target export 1:322
 route-target import 1:322
!
```

Step 2 Create loopback interfaces:

CE1 (Cisco IOS):

```
interface Loopback101
 ip vrf forwarding EC1
 ip address 172.17.3.1 255.255.255.255
!
interface Loopback102
 ip vrf forwarding EC2
 ip address 172.17.3.2 255.255.255.255
!
```

CE2 (Cisco IOS):

```
interface Loopback101
 ip vrf forwarding EC1
 ip address 172.17.4.1 255.255.255.255
!
interface Loopback102
```

```
 ip vrf forwarding EC2
 ip address 172.17.4.2 255.255.255.255
!
```

Task 3: Configure Routing Between the PE and CE Routers

Step 1 Configure a BGP process:

CE1 (Cisco IOS):

```
no router bgp <AS number>
router bgp 64511
 neighbor 192.168.101.10 remote-as 64500
 neighbor 192.168.101.10 update-source GigabitEthernet0/1
 !
 address-family ipv4
  neighbor 192.168.101.10 activate
  neighbor 192.168.101.10 next-hop-self
exit-address-family
 !
```

CE2 (Cisco IOS):

```
no router bgp <AS number>
router bgp 64511
 neighbor 192.168.102.20 remote-as 64500
 neighbor 192.168.102.20 update-source GigabitEthernet0/1
 !
 address-family ipv4
  neighbor 192.168.102.20 activate
  neighbor 192.168.102.20 next-hop-self
exit-address-family
 !
```

PE1 (Cisco IOS XR):

```
router bgp 64500
 address-family ipv4 unicast
 !
 vrf Customer_1
  rd 1:210
  neighbor 192.168.101.11
   remote-as 64511
   update-source GigabitEthernet0/0/0/0
   address-family ipv4 unicast
    route-policy pass in
    route-policy pass out
    next-hop-self
  !
```

PE2 (Cisco IOS XE):

```
router bgp 64500
 neighbor 192.168.102.21 remote-as 64511
 neighbor 192.168.102.21 update-source GigabitEthernet0/0/0
 address-family ipv4 vrf Customer_1
  neighbor 192.168.102.21 remote-as 64511
  neighbor 192.168.102.21 activate
  neighbor 192.168.102.21 next-hop-self
```

Step 2 Change the Loopback1 subnet mask:

CE1 (Cisco IOS):

```
interface Loopback1
 ip address 172.16.1.1 255.255.255.255
!
```

CE2 (Cisco IOS):

```
interface Loopback1
```



```
 ip address 172.16.2.1 255.255.255.255
!
```

Step 3 Reconfigure the routers to send labels:

CE1 (Cisco IOS):

```
router bgp 64511
 address-family ipv4
  network 172.16.1.1 mask 255.255.255.255
  neighbor 192.168.101.10 send-label
 exit-address-family
 !
```

PE1 (Cisco IOS XR):

```
router bgp 64500
 vrf Customer_1
  address-family ipv4 unicast
   allocate-label all
  address-family ipv4 labeled-unicast
   route-policy pass in
   route-policy pass out
   next-hop-self
 !
```

CE2 (Cisco IOS):

```
router bgp 64511
 address-family ipv4
  network 172.16.2.1 mask 255.255.255.255
  neighbor 192.168.102.20 send-label
 exit-address-family
 !
```

PE2 (Cisco IOS XE):

```
router bgp 64500
 address-family ipv4 vrf Customer_1
  neighbor 192.168.102.21 send-label
 !
```

Step 4 Configure a static route:

PE1 (Cisco IOS XR):

```
router static
 vrf Customer_1
  address-family ipv4 unicast
   192.168.101.11/32 GigabitEthernet0/0/0/0
  !
 !
```

Task 4: Establish a BGP Session Between Customer Carrier Routers

Step 1 Establish an IBGP session:

CE3 (Cisco IOS):

```
router bgp 64511
 bgp log-neighbor-changes
 neighbor 172.16.2.1 remote-as 64511
 neighbor 172.16.2.1 update-source Loopback1

 address-family ipv4
  no neighbor 172.16.2.1 activate
  !
 exit-address-family
 !
 address-family vpnv4
  neighbor 172.16.2.1 activate
  neighbor 172.16.2.1 send-community both
 exit-address-family
```

CE4 (Cisco IOS):

```
router bgp 64511
 bgp log-neighbor-changes
 neighbor 172.16.1.1 remote-as 64511
 neighbor 172.16.1.1 update-source Loopback1

 address-family ipv4
  no neighbor 172.16.1.1 activate
  !
 exit-address-family
 !
 address-family vpnv4
  neighbor 172.16.1.1 activate
  neighbor 172.16.1.1 send-community both
 exit-address-family
```

Step 2 Configure an AS override:

PE1 (Cisco IOS XR):

```
router bgp 64500
 vrf Customer_1
  neighbor 192.168.101.11
   address-family ipv4 unicast
    as-override
   address-family ipv4 labeled-unicast
    as-override
```

PE2 (Cisco IOS XE):

```
router bgp 64500
 address-family ipv4 vrf Customer_1
  neighbor 192.168.102.21 as-override
```

Step 3 Establish an IBGP session:

CE3 (Cisco IOS):

```
router bgp 64512

 address-family ipv4 vrf EC1
  redistribute connected
 exit-address-family
 !
 address-family ipv4 vrf EC2
  redistribute connected
 exit-address-family
!
```

CE4 (Cisco IOS):

```
router bgp 64512

 address-family ipv4 vrf EC1
  redistribute connected
 exit-address-family
 !
 address-family ipv4 vrf EC2
  redistribute connected
 exit-address-family
!
```

Lab 4-1 Answer Key: Implement Layer 2 VPN (VPWS and VPLS)

When you complete this lab activity, the device configuration and device outputs will be similar to the results shown here, with differences that are specific to your pod.

Task 1: Remove the CSC Configuration

Note	Parameters depend on the pod number.

Step 1 Remove the BGP configuration on the CE routers:

CE3, CE4 (Cisco IOS):

```
no router bgp <AS-number>
!
```

Step 2 Configure the CE-PE interfaces:

CE3, CE4 (Cisco IOS):

```
default interface GigabitEthernet0/1
!
interface GigabitEthernet0/1
 duplex full
 speed 100
!
```

PE3 (Cisco IOS XR):

```
default interface GigabitEthernet0/0/0/0
!
interface GigabitEthernet0/0/0/0
 speed 100
!
```

PE4 (Cisco IOS XE):

```
default interface GigabitEthernet0/0/0
!
interface GigabitEthernet0/0/0
 speed 100
```

Task 2: Configure EoMPLS

Note	Parameters depend on the pod number.

Step 1 Configure the IP address on the CE router:

CE3 (Cisco IOS):

```
interface GigabitEthernet0/1
 ip address 192.168.200.3 255.255.255.0
!
```

CE4 (Cisco IOS):

```
interface GigabitEthernet0/1
 ip address 192.168.200.4 255.255.255.0

!
```

Step 2 Configure Layer 2 transport on the Cisco IOS XR router:

PE3 (Cisco IOS XR):

```
interface GigabitEthernet0/0/0/0
 l2transport
!
```

Step 3 Configure EoMPLS on the Cisco IOS XR router:

PE3 (Cisco IOS XR):

```
l2vpn
 pw-class pod3-pw-class
  encapsulation mpls
  !
 !
 xconnect group pod3-group
  p2p pod3-eompls
   interface GigabitEthernet0/0/0/0
   neighbor 10.4.1.1 pw-id 34
    pw-class pod3-pw-class
!
```

Step 4 Configure Layer 2 transport on the Cisco IOS XE router:

PE4 (Cisco IOS XE):

```
pseudowire-class pw-class
 encapsulation mpls
!
interface GigabitEthernet0/0/0
 xconnect 10.3.1.1 34 pw-class pw-class
!
```

Task 3: Configure VPLS

Note	Parameters depend on the pod number.

Step 1 Configure a bridge domain:

PE1 (Cisco IOS XR):

```
l2vpn
 no xconnect group xcon_grp
 !
 bridge group BG1
  bridge-domain BD1
!
```

Step 2 Configure interface GigabitEthernet0/0/0/0:

PE1 (Cisco IOS XR):

```
l2vpn
 bridge group BG1
  bridge-domain BD1
   interface GigabitEthernet0/0/0/0
!
```

Step 3 Configure pseudowires:

PE1 (Cisco IOS XR):

```
l2vpn
 bridge group BG1
  bridge-domain BD1
    neighbor 10.2.1.1 pw-id 12
     pw-class pod1-pw-class
    !
    neighbor 10.3.1.1 pw-id 13
     pw-class pod1-pw-class
!
```


Task 4: Use a VFI

Note	Parameters depend on the pod number.

Step 1 Configure a VFI:

PE1 (Cisco IOS XR):

```
l2vpn
 bridge group BG1
  bridge-domain BD1
   vfi vfi1
!
```

Step 2 Move the pseudowires:

PE1 (Cisco IOS XR):

```
l2vpn
 bridge group BG1
  bridge-domain BD1
   no neighbor 10.2.1.1 pw-id 12
   no neighbor 10.3.1.1 pw-id 13
   exit
  vfi vfi1
    neighbor 10.2.1.1 pw-id 10
     pw-class pod1-pw-class
    !
    neighbor 10.3.1.1 pw-id 13
     pw-class pod1-pw-class
!
```

Lab Topology

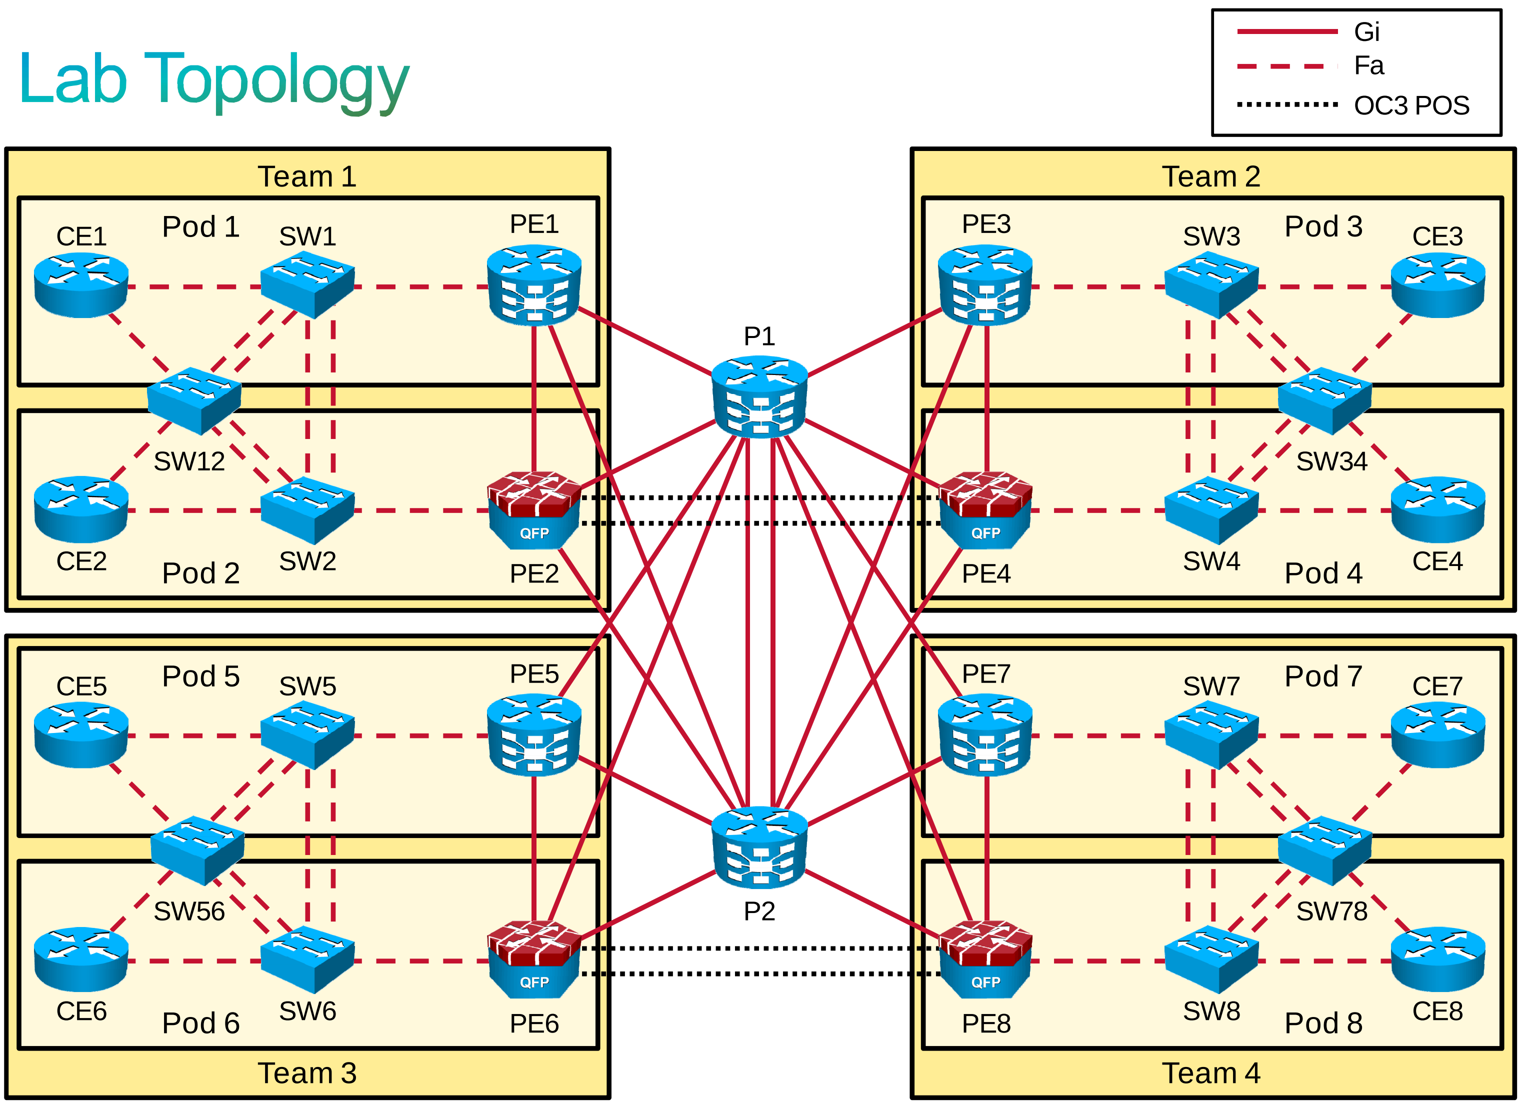

SPEDGE v1.0-4

Lab Interface Identification—Team View

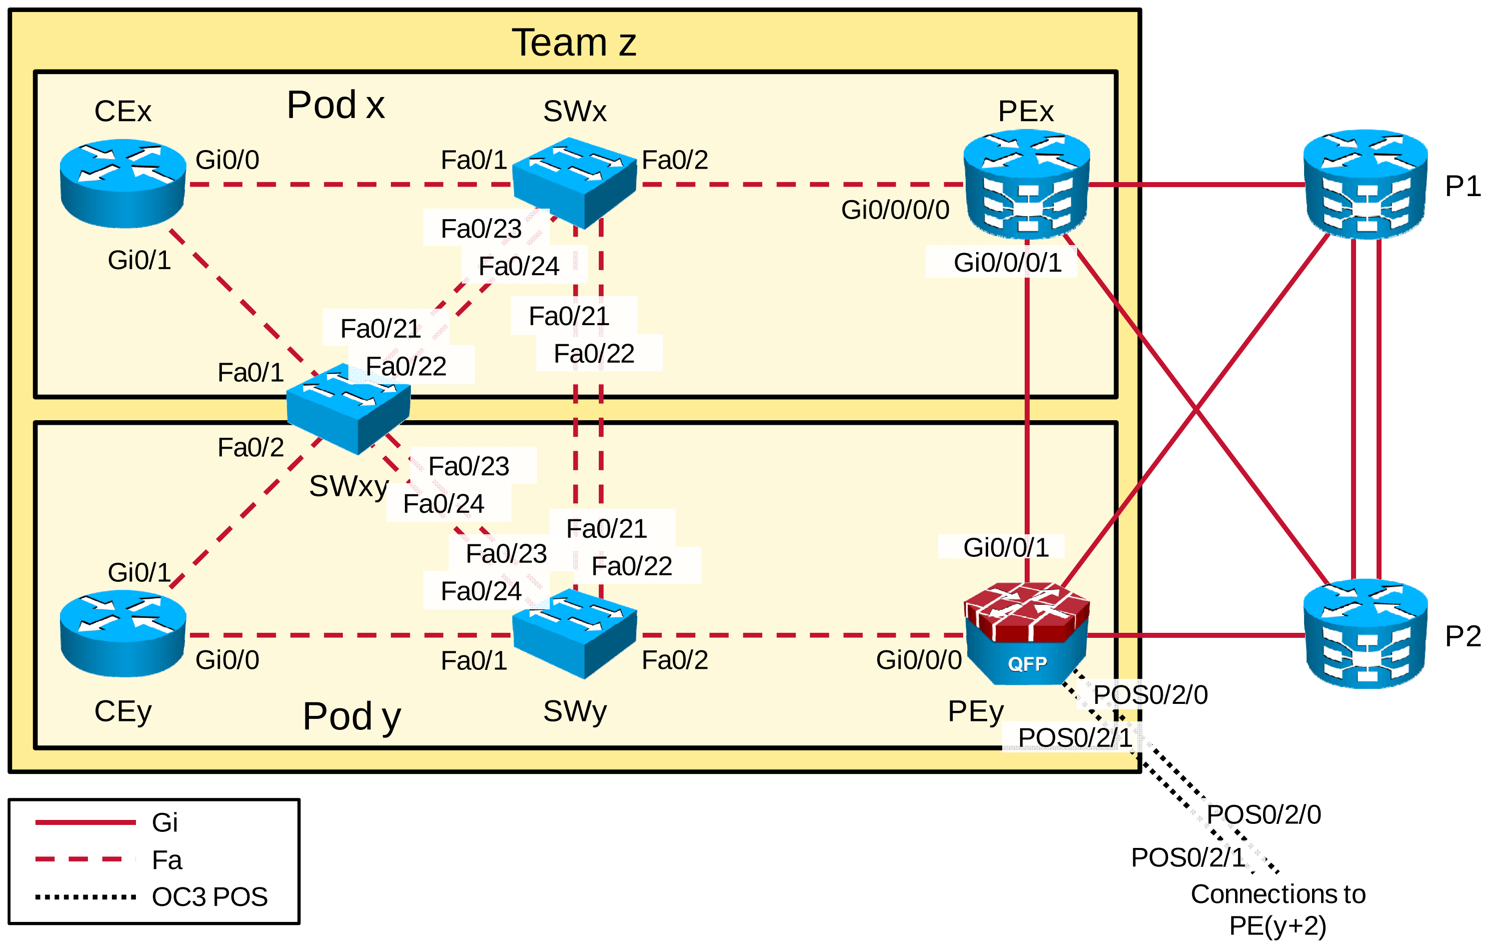

SPEDGE v1.0-5

Lab IP Addressing—Team View

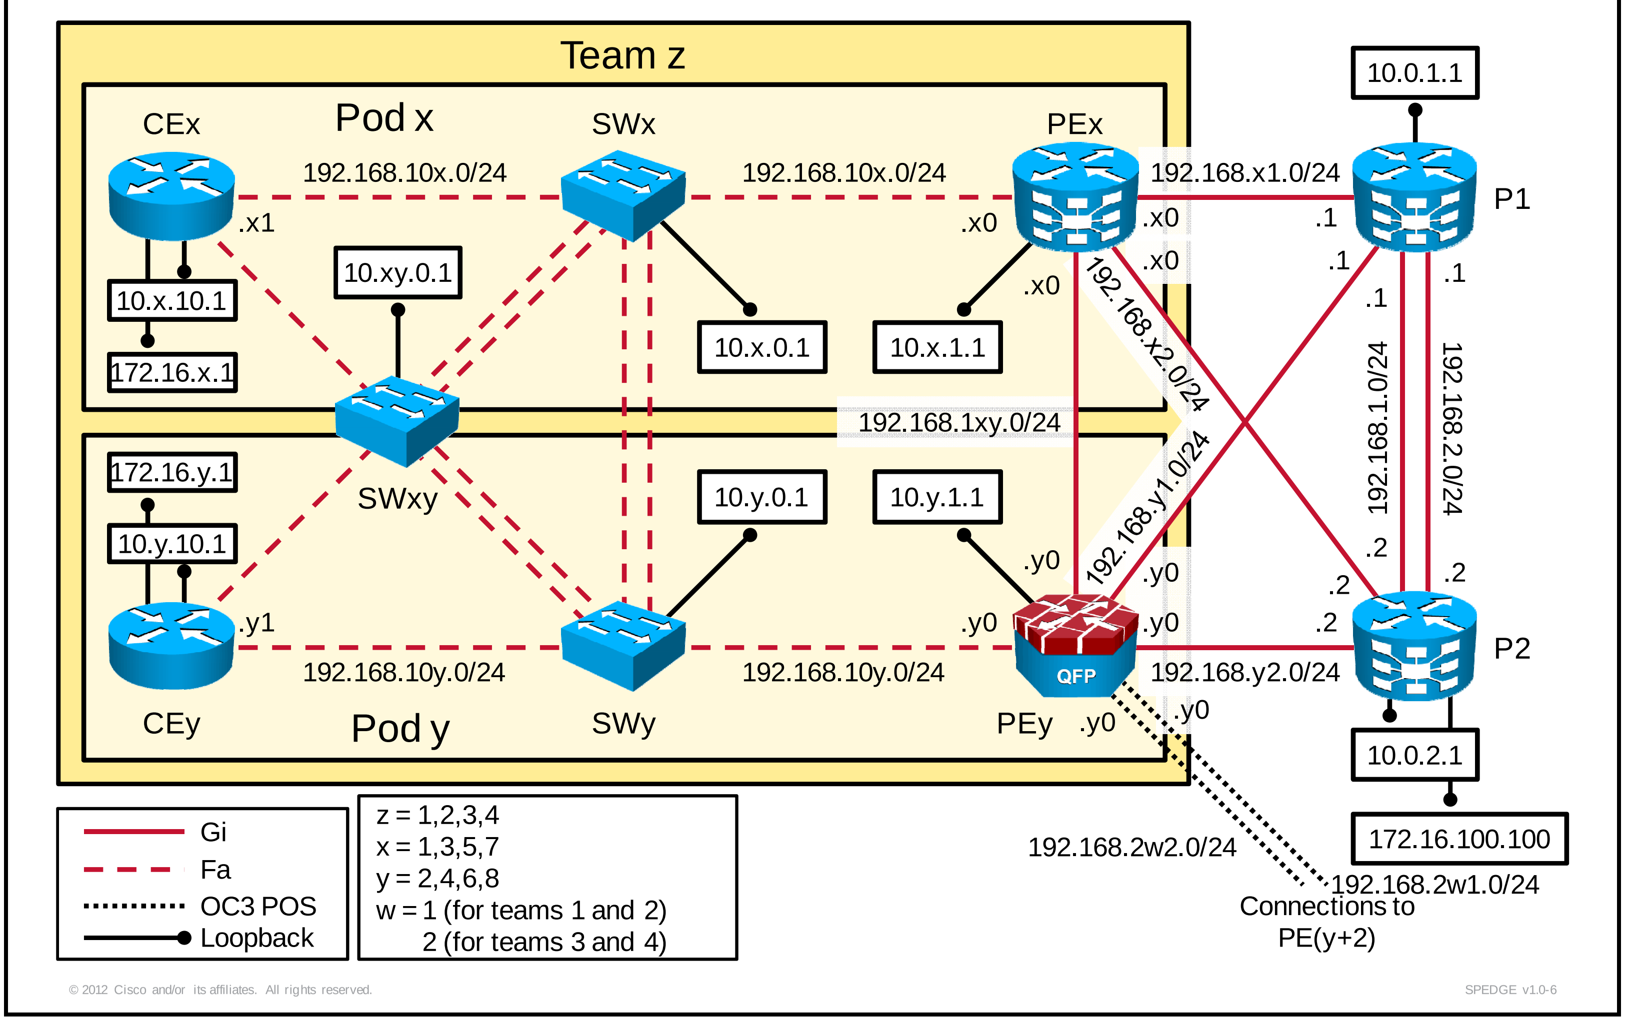

Management IP Addresses

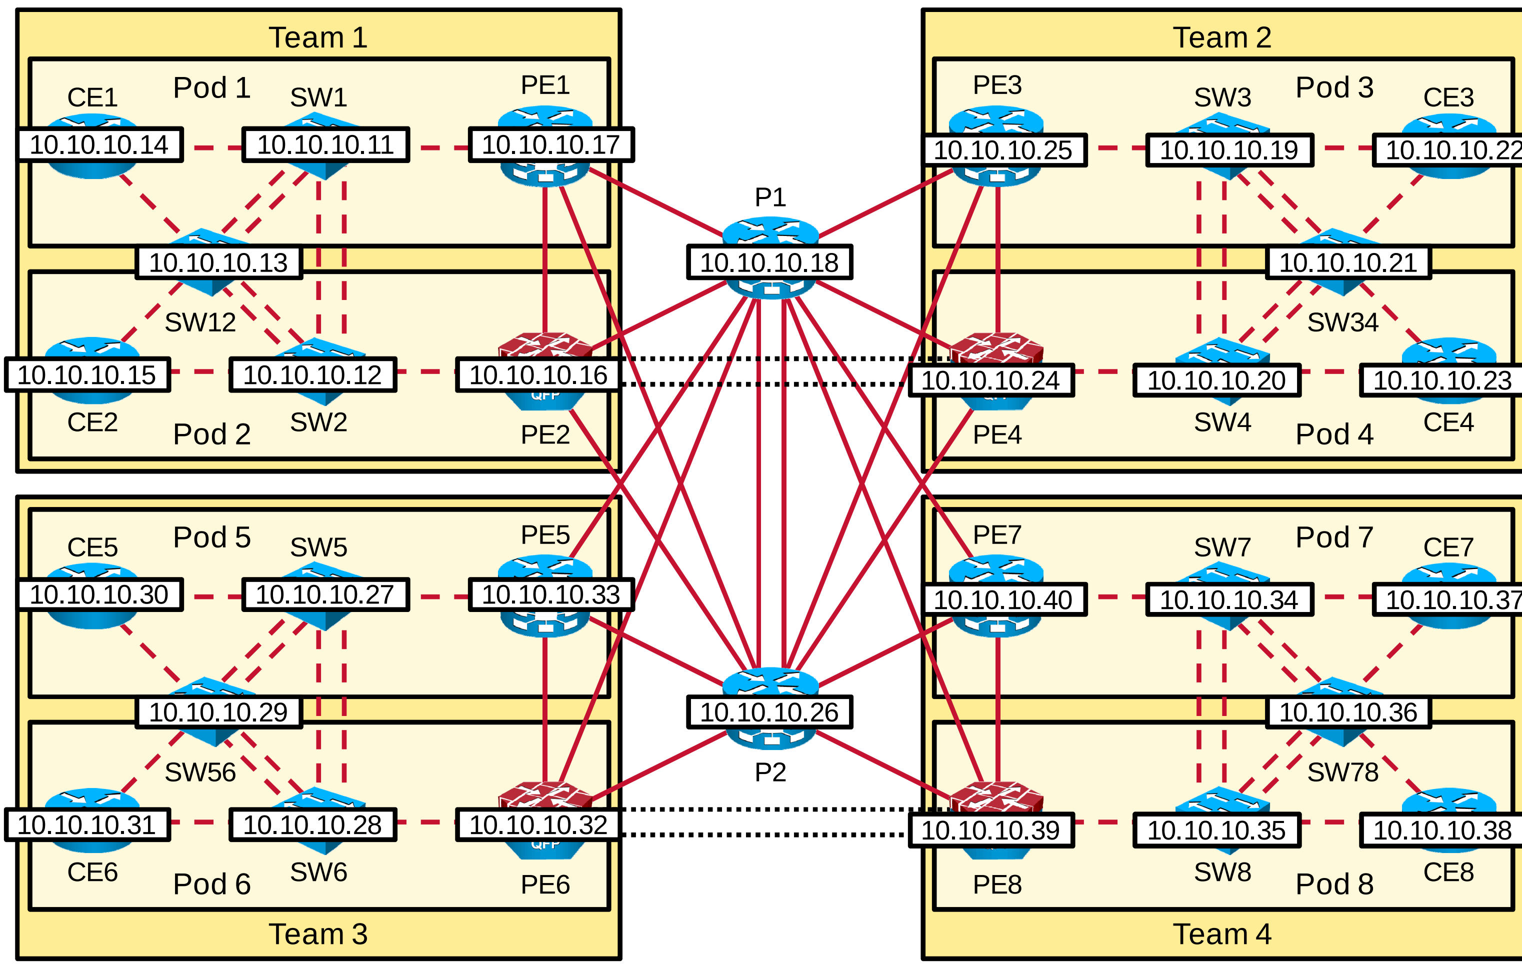

SPEDGE v1.0-7