



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.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

DISCLAIMER WARRANTY: THIS CONTENT IS BEING PROVIDED "AS IS" AND AS SUCH MAY INCLUDE TYPOGRAPHICAL, GRAPHICS, OR FORMATTING ERRORS. CISCO MAKES AND YOU RECEIVE NO WARRANTIES IN CONNECTION WITH THE CONTENT PROVIDED HEREUNDER, EXPRESS, IMPLIED, STATUTORY OR IN ANY OTHER PROVISION OF THIS CONTENT OR COMMUNICATION BETWEEN CISCO AND YOU. CISCO SPECIFICALLY DISCLAIMS ALL IMPLIED WARRANTIES, INCLUDING WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT AND FITNESS FOR A PARTICULAR PURPOSE, OR ARISING FROM A COURSE OF DEALING, USAGE OR TRADE PRACTICE. This learning product may contain early release content, and while Cisco believes it to be accurate, it falls subject to the disclaimer above.

Copyright Date:

© 2014 Cisco Systems, Inc.

Table of Contents


Lab Introduction	1
Lab Topology Diagram (Backbone Pod View)	1
Lab Topology Diagram (Student Pod View)	2
Global Job Aids	3
Pod Access Information	3
Device Information	4
IP Addressing	6
Hardware Lab 1: Implement OSPF Routing	9
Visual Objective	10
OSPF Areas	10
Required Resources	11
Command List	12
Task 1: Enabling OSPF on the Routers	14
Task 2: Influencing OSPF DR and BDR Election	27
Task 3: Influencing OSPF Route Selection by Changing the OSPF Link Cost	28
Task 4: Configuring OSPF Authentication	31
Task 5: Configuring OSPF Virtual Links	33
Hardware Lab 2: Implement OSPF Special Area Types	37
Visual Objective	38
OSPF Areas	38
Required Resources	39
Command List	40
Task 1: Enabling OSPF Summarization	41
Task 2: Configuring the OSPF Stub Area	43
Hardware Lab 3: Implement Integrated IS-IS Routing	49
Visual Objective	50
Required Resources	50
Command List	51
Task 1: Enabling Integrated IS-IS on the Routers	53
Task 2: IS-IS Route Summarization	59
Task 3: Enabling IPv6 IS-IS Routing	61
Hardware Lab 4: Implement Basic BGP Routing	65
Visual Objective	66
Required Resources	66
Command List	67
Task 1: Configuring External BGP	70
Task 2: Configuring Internal BGP	73
Hardware Lab 5: Implement Route Redistribution	77
Visual Objective	78
Required Resources	78
Command List	79
Task 1: Redistributing between OSPF and IS-IS	81
Task 2: Preventing Potential Routing Loops	86
Task 3: Modifying Administrative Distance	89
Task 4: One-Way Redistribution	93
Hardware Lab 6: Influence BGP Route Selection	95
Visual Objective	96
Required Resources	96

Command List.....	97
Task 1: Setting the Second Link between the CE and PE Routers.....	99
Task 2: Configuring BGP Weight and Local Preference.....	100
Task 3: Configuring BGP Multiexit Discriminator.....	103
Task 4: Configuring the BGP Community.....	106
Answer Key	117
Hardware Lab 1: Implement OSPF Routing	117
Hardware Lab 2: Implement OSPF Special Area Types.....	126
Hardware Lab 3: Implement Integrated IS-IS Routing.....	130
Hardware Lab 4: Implement Basic BGP Routing.....	134
Hardware Lab 5: Implement Route Redistribution	138
Hardware Lab 6: Influence BGP Route Selection.....	142


Lab Introduction

Follow this lab guide to practice what you have learned in the course.

Lab Topology Diagram (Backbone Pod View)



Lab Topology Diagram (Student Pod View)



Global Job Aids

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

Pod Access Information

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


Parameter	Default Value	Value
Team number	z = 1 or 2	
Pod number	x = 1, 3 or y = 2, 4	
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	IronMan	
Pod CE, SW, and PE privileged level password	cisco	

Device Information

This lab topology consists of two (2) teams and four (4) 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. Two 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 (pods 1 or 3) will work on the PE router running Cisco IOS XR Software, and the second pod within the same team (pods 2 or 4) will work on the PE router 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 following topology:




Device Roles and Loopback IP Addresses

Device Name	Device Role	Lo0 IPv4 Address	Lo0 IPv6 Address
CEx	Cisco 2900 pod router	10.x.10.1/32	2001:db8:10:x:10::1/128
CEy		10.y.10.1/32	2001:db8:10:y:10::1/128
PEx	Cisco ASR 9000 or Cisco ASR 1000 pod router	10.x.1.1/32	2001:db8:10:x:1::1/128
PEy		10.y.1.1/32	2001:db8:10:y:1::1/128
SWx	Cisco ME340x pod switch	10.x.0.1/32	2001:db8:10:x:0::1/128
SWy		10.y.0.1/32	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


Device Name	Device Role	Lo0 IPv4 Address	Lo0 IPv6 Address
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 following figure illustrates the interface identification that is used in this lab setup.




IP Addressing

The following figure illustrates the IP addressing scheme that is 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 access the IP addresses within your pod (x is for odd number pods 1 and 3; y is for even number pods 2 and 4). Replace the xy (where x < y) with numbers of the pods within the same team (for example, 12 or 34) to access IP addresses on the link between those pods.

Pod IP Addressing

Device	Interface	IPv4 Address	IPv6 Address
CEx	Gi0/0	192.168.10x.x1/24	2001:db8:192:168:10x::x1/80
CEy	Gi0/0	192.168.10y.y1/24	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
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
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

Hardware Lab 1: Implement OSPF Routing

Activity Objective

In this lab activity, you will configure OSPF routing by enabling OSPF Area 0 to run in the backbone and other nonbackbone areas to run in your pod.


Note	Students from two different pods are working as 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 running Cisco IOS XR Software, and the second pod in the same team will work on the PE router running Cisco IOS XE Software. Students in the same team should coordinate their lab activity.
-------------	---

In the lab activity, you will work on different Cisco routers 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 routers with OSPFv2 and OSPFv3 routing protocols and stable OSPF router IDs
- Influence OSPF DR and BDR election on a LAN
- Influence OSPF route selection by changing the OSPF link cost
- Configure OSPF authentication
- Configure OSPF virtual links to support OSPF areas not directly connected to Area 0

Visual Objective

The figure illustrates what you will accomplish in this activity.



OSPF Areas

This section includes a table with OSPF areas that are used in the lab. Replace the x or y with your pod number to access the OSPF areas that are used in the lab.

Link	OSPF Area
P1 – P2	0
P1 and P2 Loopback0	
PEx – P1, PEy – P1	
PEx – P2, PEy – P2	
PEx – PEy	
PEx Loopback0, PEy Loopback0	x, y
CEx – PEX, CEy – PEy	1x, 1y
CEx Loopback0, CEy Loopback0	
CEx Loopback1x (1x.0.0.1/24), CEy Loopback1y (1y.0.0.1/24)	x00, y00

Required Resources

These are the resources and equipment that 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 the interface
area <i>area</i> virtual-link <i>IP-address</i>	(Router) Enables the OSPF virtual link
clear ip ospf process	Clears OSPF processes on the router
configure terminal	Enters configuration mode
interface <i>interface</i>	Enters interface configuration mode
ip ospf authentication message-digest	Enables OSPF MD5 authentication on the interface
ip ospf cost <i>cost</i>	Sets the OSPF link cost on the interface
ip ospf message-digest-key <i>key-ID</i> md5 <i>key</i>	Sets the OSPF MD5 key on the interface
ip ospf priority <i>priority</i>	Sets the OSPF interface priority for DR and BDR election
ip ipv6 address <i>ip_address subnet_mask</i>	Sets an IPv4 or IPv6 address and the subnet mask on the interface
ipv6 enable	Enables IPv6 support on the interface
ipv6 ospf <i>process-id</i> <i>area area</i>	Enables OSPFv3 routing on the interface
ipv6 unicast-routing	Globally enables IPv6 unicast routing
network <i>prefix wildcard_mask</i> <i>area area</i>	(Router) Enables OSPF routing on the network and places the network into the OSPF area
ping <i>dest_IP</i> <i>source source_IP</i>	Verifies connectivity between the source IP and destination IP
router ospf <i>process-id</i>	Enables the OSPFv2 process on the router
router-id <i>IP-address</i>	(Router) Sets the OSPF router ID
show ip interface brief	Displays the interface status and IPv4 addresses configured
show ip ospf	Displays information related to the OSPF routing protocol running on the router

Command	Description
show ip ospf interface	Displays the OSPF interface information
show ip ospf virtual-links	Displays the OSPF virtual link information
show ip ipv6 ospf database	Displays the content of the OSPF database
show ip ipv6 ospf neighbors	Displays OSPF neighbor information
show ip ipv6 protocols	Displays IPv4 or IPv6 protocols running on the router
show ip ipv6 route	Displays the current routes in the routing table
show running-config	Displays running configuration

Cisco IOS XR Commands

Command	Description
[no] shutdown	Enables or disables the interface
area <i>area</i>	(Router) Creates the OSPF area on the router
authentication message-digest	(Router) Enables OSPF MD5 authentication on the interface
clear ospf process	Clears OSPF processes on the router
commit	Commits changes to the running configuration
configure terminal	Enters configuration mode
cost <i>cost</i>	(Router) Sets the OSPF link cost on the interface
interface <i>interface</i>	(Global) Enters interface configuration mode
interface <i>interface</i>	(Router) Defines the interfaces on which the OSPF protocol runs
ipv4 ipv6 address <i>ip_address</i>/<i>n</i>	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
message-digest-key <i>key-ID</i> md5 <i>key</i>	(Router) Sets the OSPF MD5 key on the interface
ping <i>dest_IP</i> source [<i>source_IP</i>]	Verifies connectivity between the source IP and destination IP (IPv4 and IPv6)

Command	Description
priority <i>priority</i>	(Router) Sets the OSPF interface priority for DR and BDR election
router ospf ospfv3 <i>process-ID</i>	Creates an OSPFv2 or OSPFv3 process
router-id <i>IP-address</i>	(Router) Sets the OSPF router ID
show ipv4 interface brief	Displays the interface status and IPv4 addresses configured
show ospf	Displays information related to the OSPF routing protocol running on the router
show ospf interface	Displays OSPF interface information
show ospf ospfv3 database	Displays the content of the OSPF database
show ospf ospfv3 neighbors	Displays OSPF neighbor information
show protocols	Displays protocols running on the router
show route	Displays the current routes in the routing table
show running-config	Displays running configuration
traceroute <i>IP-address</i>	Traces the IP address
virtual-link <i>IP-address</i>	(Router) Enables the OSPF virtual link

Task 1: Enabling OSPF on the Routers

In this task, you will enable OSPF version 2 and OSPF version 3 routing protocols, and you will configure stable OSPF router IDs.

Note	Outputs in this task are taken from team 2 with pods 3 and 4.
-------------	---

Activity Procedure

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

Step 1 On the CE router, enable the Gi0/0 and Lo0 interfaces and configure IP addresses. On the PE router, enable the Gi0/0/0/0, Gi0/0/0/1, Gi0/0/0/2, Gi0/0/0/3, and Loopback0 interfaces and configure IP addresses. Find the correct IP addresses in the job aids.

On the CE routers, there should first be Gigabit Ethernet and Loopback interfaces up and running with assigned IP addresses.

```
CE3# show ip interface brief | include GigabitEthernet0/0|Loopback0
GigabitEthernet0/0      192.168.103.31 YES manual up
Loopback0               10.3.10.1     YES manual up
```

```
CE4# show ip interface brief | include GigabitEthernet0/0|Loopback0
GigabitEthernet0/0      192.168.104.41 YES manual up
Loopback0               10.4.10.1     YES manual up
```

On the PE routers, there should be four Gigabit Ethernet and Loopback interfaces up and running with assigned IP addresses.

```
RP/0/RSP0/CPU0:PE3# show ipv4 interface brief | include Up
Wed May 14 11:59:13.620 UTC
Loopback0               10.3.1.1      Up
MgmtEth0/RSP0/CPU0/0    10.10.10.25   Up
GigabitEthernet0/0/0/0   192.168.103.30 Up
GigabitEthernet0/0/0/1   192.168.134.30 Up
GigabitEthernet0/0/0/2   192.168.31.30 Up
GigabitEthernet0/0/0/3   192.168.32.30 Up
```

```
PE4# show ip interface brief | include up
GigabitEthernet0/0/0     192.168.104.40 YES manual up
GigabitEthernet0/0/1     192.168.134.40 YES manual up
GigabitEthernet0/0/2     192.168.41.40  YES manual up
GigabitEthernet0/0/3     192.168.42.40  YES manual up
GigabitEthernet0         10.10.10.24    YES NVRAM up
Loopback0                10.4.1.1       YES manual up
```

Step 2 You should be able to successfully test IP connectivity between the CE and PE routers in your pod, between the PE routers in your team, and between the PE router in your team and core routers P1 and P2.

The ping is successful from the PE to the CE router:

```
RP/0/RSP0/CPU0:PE3# ping 192.168.103.31
Wed May 14 12:39:29.211 UTC
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.103.31, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/10 ms
```

The ping is successful between the PE routers in your team:

```
RP/0/RSP0/CPU0:PE3# ping 192.168.134.40
Wed May 14 12:40:02.137 UTC
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.134.40, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/3 ms
```

The ping is successful between the PE router in your team and core routers P1 and P2:

```
RP/0/RSP0/CPU0:PE3# ping 192.168.31.1
Wed May 14 12:40:53.250 UTC
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.31.1, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/12 ms

RP/0/RSP0/CPU0:PE3# ping 192.168.32.2
Wed May 14 12:41:28.582 UTC
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.32.2, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/3 ms
```

In the following steps, you will configure basic OSPF routing in your team routers.

Step 3 Determine the OSPF area for every link in your team. Use the visual objective and OSPF Areas sections to complete this step:

Link (Interface)	OSPF Area
PE _x (Gi0/0/0/2) – P1	
PE _x (Gi0/0/0/3) – P2	
PE _x (Gi0/0/0/1) – PE _y (Gi0/0/0/1)	
PE _x (Gi0/0/0/0) – CE _x (Gi0/0)	
PE _x (Loopback0)	
CE _x (Loopback0)	
PE _y (Gi0/0/0/2) – P1	
PE _y (Gi0/0/0/3) – P2	
PE _y (Gi0/0/0/0) – CE _y (Gi0/0)	
PE _y (Loopback0)	
CE _y (Loopback0)	

Step 4 On the CE and PE routers in your pod, enable OSPF routing and assign interfaces to the areas that are determined in the previous step. Use OSPF process ID 1.

Step 5 You should be able to successfully test IP connectivity between the CE routers in your team. The ping is successful from the CE router in your pod and CE router in the neighbor pod (same team).

```
CE3# ping 10.4.10.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.4.10.1, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

Step 6 Verify which router ID is assigned to the OSPF routing process on the CE and PE routers in your team.

The OSPF router ID is on the CE and PE routers:

```
CE3# show ip ospf | include ID
Routing Process "ospf 1" with ID 10.3.10.1
```

```
RP/0/RSP0/CPU0:PE3# show ospf | include ID
Thu May 15 09:44:54.272 UTC
Routing Process "ospf 1" with ID 10.3.1.1
```

Write the results into the following table:

Router	OSPF Router ID
PE _x	
CE _x	
PE _y	
CE _y	

Step 7 On the CE and PE routers, set the OSPF router IDs to the following IP addresses:

Router	OSPF Router ID
PE _x	x.x.x.x
CE _x	x0.x0.x0.x0
PE _y	y.y.y.y
CE _y	y0.y0.y0.y0

Make sure that the new OSPF router IDs will take effect.

- Step 8** On the CE and PE routers in your pod, enable IPv6 routing. On the CE router, enable IPv6 on the Gi0/0 and Lo0 interfaces and configure the IPv6 address on the Lo0 interface only. On the PE router, enable IPv6 on the Gi0/0/0/0, Gi0/0/0/1, Gi0/0/0/2, Gi0/0/0/3, and Lo0 interfaces and configure the IP address on the Lo0 interface only. Find the correct IPv6 addresses in the job aids.
- Step 9** On the CE and PE routers in your pod, enable OSPF version 3 routing and assign interfaces to the areas as used for OSPF version 2.

Verification

You have completed this task when you attain these results:

- Step 1** On the CE router in your pod, verify the IP protocols running, OSPF neighbors, and OSPF database. There should be one OSPF neighbor. The OSPF database should show LSAs from Area 1x or 1y, where x or y is your pod number. Outputs should be similar to the following:

```
CE3# show ip protocols
```

```
*** IP Routing is NSF aware ***
```

```
Routing Protocol is "ospf 1"
```

```
Outgoing update filter list for all interfaces is not set
```

```
Incoming update filter list for all interfaces is not set
```

```
Router ID 30.30.30.30
```

```
Number of areas in this router is 1. 1 normal 0 stub 0 nssa
```

```
Maximum path: 4
```

```
Routing for Networks:
```

```
10.3.10.1 0.0.0.0 area 13
```

```
192.168.103.31 0.0.0.0 area 13
```

```
Routing Information Sources:
```

```
Gateway Distance Last Update
```

```
3.3.3.3 110 19:35:11
```

```
10.3.1.1 110 19:41:58
```

```
Distance: (default is 110)
```

```
CE3# show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
3.3.3.3	1	FULL/BDR	00:00:34	192.168.103.30	GigabitEthernet0/0

```
CE3# show ip ospf database
```

```
OSPF Router with ID (30.30.30.30) (Process ID 1)
```

```
Router Link States (Area 13)
```

Link ID	ADV Router	Age	Seq#	Checksum	Link count
3.3.3.3	3.3.3.3	1587	0x80000028	0x0036DD	1
30.30.30.30	30.30.30.30	1651	0x80000027	0x00F919	2

```
Net Link States (Area 13)
```

Link ID	ADV Router	Age	Seq#	Checksum
192.168.103.31	30.30.30.30	1651	0x80000025	0x000C1D

```
Summary Net Link States (Area 13)
```

Link ID	ADV Router	Age	Seq#	Checksum
10.0.1.1	3.3.3.3	569	0x80000025	0x00639B
10.0.2.1	3.3.3.3	569	0x80000025	0x0058A5
10.1.1.1	3.3.3.3	569	0x80000025	0x00619B
10.2.1.1	3.3.3.3	569	0x80000025	0x0055A6
10.3.1.1	3.3.3.3	569	0x80000025	0x0035C7
10.4.1.1	3.3.3.3	316	0x80000029	0x002BCB
10.4.10.1	3.3.3.3	316	0x80000024	0x00DB16
192.168.1.0	3.3.3.3	569	0x80000025	0x003E62
192.168.2.0	3.3.3.3	569	0x80000025	0x008D09
192.168.11.0	3.3.3.3	569	0x80000025	0x00CFC6
192.168.12.0	3.3.3.3	569	0x80000025	0x002962
192.168.21.0	3.3.3.3	569	0x80000025	0x00C5BC
192.168.22.0	3.3.3.3	569	0x80000025	0x005635
192.168.31.0	3.3.3.3	569	0x80000025	0x00E89A
192.168.32.0	3.3.3.3	569	0x80000027	0x00D9A6
192.168.41.0	3.3.3.3	316	0x80000029	0x007CF7
192.168.42.0	3.3.3.3	569	0x80000025	0x0079FD
192.168.104.0	3.3.3.3	316	0x80000024	0x00CE6B
192.168.134.0	3.3.3.3	316	0x80000028	0x0071A7

Step 2 On the PE router in your pod, verify the IP protocols running, OSPF neighbors, and OSPF database. There should be four OSPF neighbors. The OSPF database should show LSAs from Areas 0, x, or y, and lx or ly, where x or y is your pod number. Outputs should be similar to the following:

```
RP/0/RSP0/CPU0:PE3# show protocols
Fri May 16 07:49:05.610 UTC
```

```
Routing Protocol OSPF 1
  Router Id: 3.3.3.3
  Distance: 110
  Non-Stop Forwarding: Disabled
  Redistribution:
    None
  Area 0
    GigabitEthernet0/0/0/1
    GigabitEthernet0/0/0/2
    GigabitEthernet0/0/0/3
  Area 3
    Loopback0
  Area 13
    GigabitEthernet0/0/0/0
```

```
RP/0/RSP0/CPU0:PE3# show ospf neighbor
Fri May 16 07:49:14.729 UTC
```

* Indicates MADJ interface

Neighbors for OSPF 1

Neighbor ID	Pri	State	Dead Time	Address	Interface
4.4.4.4	1	FULL/BDR	00:00:35	192.168.134.40	GigabitEthernet0/0/0/1
Neighbor is up for 19:39:00					
0.0.0.1	1	FULL/DR	00:00:39	192.168.31.1	GigabitEthernet0/0/0/2
Neighbor is up for 19:43:27					
0.0.0.2	1	FULL/DR	00:00:37	192.168.32.2	GigabitEthernet0/0/0/3
Neighbor is up for 19:43:24					
30.30.30.30	1	FULL/DR	00:00:34	192.168.103.31	GigabitEthernet0/0/0/0
Neighbor is up for 17:15:42					

Total neighbor count: 4

```
RP/0/RSP0/CPU0:PE3# show ospf database
Fri May 16 07:49:23.265 UTC
```

OSPF Router with ID (3.3.3.3) (Process ID 1)

Router Link States (Area 0)

Link ID	ADV Router	Age	Seq#	Checksum	Link count
0.0.0.1	0.0.0.1	326	0x800002e3	0x004a73	7
0.0.0.2	0.0.0.2	421	0x800002e4	0x00248d	7
1.1.1.1	1.1.1.1	1148	0x800002d1	0x00f23a	3
2.2.2.2	2.2.2.2	995	0x8000027b	0x00958c	3
3.3.3.3	3.3.3.3	497	0x8000002b	0x003917	3
4.4.4.4	4.4.4.4	261	0x80000026	0x006e98	3

Net Link States (Area 0)

Link ID	ADV Router	Age	Seq#	Checksum
192.168.1.2	0.0.0.2	421	0x800002c6	0x0057a9
192.168.2.2	0.0.0.2	421	0x800002c6	0x004cb3
192.168.11.10	1.1.1.1	1148	0x800002c7	0x009a4f
192.168.12.2	0.0.0.2	421	0x800002c4	0x0002f2
192.168.21.20	2.2.2.2	995	0x80000279	0x005cbf
192.168.22.20	2.2.2.2	995	0x800002c7	0x00c209
192.168.31.1	0.0.0.1	585	0x8000002c	0x00d1a5
192.168.32.2	0.0.0.2	920	0x8000002c	0x00c0b2
192.168.41.1	0.0.0.1	326	0x8000002a	0x0099d1
192.168.42.2	0.0.0.2	421	0x8000002a	0x0088de
192.168.134.30	3.3.3.3	497	0x80000024	0x008759

Summary Net Link States (Area 0)

Link ID	ADV Router	Age	Seq#	Checksum
10.3.1.1	3.3.3.3	749	0x80000025	0x0035c7
10.3.10.1	3.3.3.3	1767	0x8000001f	0x00e711
10.4.1.1	4.4.4.4	261	0x80000024	0x000deb
10.4.10.1	4.4.4.4	14	0x80000024	0x00b33b
192.168.103.0	3.3.3.3	1767	0x80000029	0x00c571
192.168.104.0	4.4.4.4	14	0x80000024	0x00a690

Type-10 Opaque Link Area Link States (Area 0)

Link ID	ADV Router	Age	Seq#	Checksum	Opaque ID
1.0.0.0	0.0.0.1	1104	0x800002c7	0x001448	0
1.0.0.0	0.0.0.2	1415	0x800002a6	0x005922	0
1.0.0.9	0.0.0.2	1415	0x800002a7	0x0058fb	9
1.0.0.10	0.0.0.2	1415	0x800002a8	0x0093aa	10
1.0.0.11	0.0.0.2	1415	0x800002a8	0x0025ef	11
1.0.0.12	0.0.0.1	326	0x800002c6	0x00a90f	12
1.0.0.12	0.0.0.2	1160	0x80000028	0x000e60	12
1.0.0.13	0.0.0.1	326	0x800002c7	0x00e4bd	13
1.0.0.13	0.0.0.2	1415	0x80000277	0x003c08	13
1.0.0.14	0.0.0.1	1104	0x800002c7	0x003a5f	14
1.0.0.14	0.0.0.2	1160	0x80000028	0x003549	14
1.0.0.15	0.0.0.1	1600	0x80000028	0x00d4c1	15
1.0.0.16	0.0.0.1	1104	0x80000277	0x008e90	16
1.0.0.17	0.0.0.1	1600	0x80000028	0x0071e8	17

Router Link States (Area 3)

Link ID	ADV Router	Age	Seq#	Checksum	Link count
3.3.3.3	3.3.3.3	750	0x80000024	0x0044a2	1

Summary Net Link States (Area 3)

Link ID	ADV Router	Age	Seq#	Checksum
10.0.1.1	3.3.3.3	749	0x80000024	0x00659a
10.0.2.1	3.3.3.3	749	0x80000025	0x0058a5
10.1.1.1	3.3.3.3	749	0x80000024	0x00639a
10.2.1.1	3.3.3.3	749	0x80000025	0x0055a6
10.3.10.1	3.3.3.3	1767	0x8000001f	0x00e711
10.4.1.1	3.3.3.3	497	0x80000025	0x0033c7
10.4.10.1	3.3.3.3	497	0x80000024	0x00db16
192.168.1.0	3.3.3.3	749	0x80000024	0x004061
192.168.2.0	3.3.3.3	749	0x80000024	0x008f08
192.168.11.0	3.3.3.3	749	0x80000024	0x00d1e5
192.168.12.0	3.3.3.3	749	0x80000024	0x002b61
192.168.21.0	3.3.3.3	749	0x80000025	0x00c5bc
192.168.22.0	3.3.3.3	749	0x80000025	0x005635
192.168.31.0	3.3.3.3	749	0x80000024	0x00ea99
192.168.32.0	3.3.3.3	749	0x80000024	0x00dfa3
192.168.41.0	3.3.3.3	497	0x80000029	0x007cf7
192.168.42.0	3.3.3.3	749	0x80000025	0x0079fd
192.168.103.0	3.3.3.3	1767	0x80000025	0x00cd6d
192.168.104.0	3.3.3.3	497	0x80000024	0x00ce6b
192.168.134.0	3.3.3.3	497	0x80000028	0x0071a7

Router Link States (Area 13)

Link ID	ADV Router	Age	Seq#	Checksum	Link count
3.3.3.3	3.3.3.3	1767	0x80000028	0x0036dd	1
30.30.30.30	30.30.30.30	1833	0x80000027	0x00f919	2

Net Link States (Area 13)

Link ID	ADV Router	Age	Seq#	Checksum
192.168.103.31	30.30.30.30	1833	0x80000025	0x000c1d

Summary Net Link States (Area 13)

Link ID	ADV Router	Age	Seq#	Checksum
10.0.1.1	3.3.3.3	749	0x80000025	0x00639b
10.0.2.1	3.3.3.3	749	0x80000025	0x0058a5
10.1.1.1	3.3.3.3	749	0x80000025	0x00619b
10.2.1.1	3.3.3.3	749	0x80000025	0x0055a6
10.3.1.1	3.3.3.3	749	0x80000025	0x0035c7
10.4.1.1	3.3.3.3	497	0x80000029	0x002bcb
10.4.10.1	3.3.3.3	497	0x80000024	0x00db16
192.168.1.0	3.3.3.3	749	0x80000025	0x003e62
192.168.2.0	3.3.3.3	749	0x80000025	0x008d09
192.168.11.0	3.3.3.3	749	0x80000025	0x00cfc6
192.168.12.0	3.3.3.3	749	0x80000025	0x002962
192.168.21.0	3.3.3.3	749	0x80000025	0x00c5bc
192.168.22.0	3.3.3.3	749	0x80000025	0x005635
192.168.31.0	3.3.3.3	749	0x80000025	0x00e89a
192.168.32.0	3.3.3.3	749	0x80000027	0x00d9a6
192.168.41.0	3.3.3.3	497	0x80000029	0x007cf7
192.168.42.0	3.3.3.3	749	0x80000025	0x0079fd
192.168.104.0	3.3.3.3	497	0x80000024	0x00ce6b
192.168.134.0	3.3.3.3	497	0x80000028	0x0071a7

Step 3 On the CE and PE routers, verify that OSPF routes are present in the IP routing table.

Output on the CE router should be similar to the following:

```
CE3# show ip route ospf
```

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

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

```
10.0.0.0/32 is subnetted, 8 subnets  
O IA 10.0.1.1 [110/3] via 192.168.103.30, 19:45:41, GigabitEthernet0/0  
O IA 10.0.2.1 [110/3] via 192.168.103.30, 19:45:41, GigabitEthernet0/0  
O IA 10.1.1.1 [110/4] via 192.168.103.30, 19:45:41, GigabitEthernet0/0  
O IA 10.2.1.1 [110/4] via 192.168.103.30, 19:45:41, GigabitEthernet0/0  
O IA 10.3.1.1 [110/2] via 192.168.103.30, 19:45:41, GigabitEthernet0/0  
O IA 10.4.1.1 [110/3] via 192.168.103.30, 19:41:24, GigabitEthernet0/0  
O IA 10.4.10.1 [110/4] via 192.168.103.30, 19:40:54, GigabitEthernet0/0  
O IA 192.168.1.0/24 [110/3] via 192.168.103.30, 19:45:41, GigabitEthernet0/0  
O IA 192.168.2.0/24 [110/12] via 192.168.103.30, 19:45:41, GigabitEthernet0/0  
O IA 192.168.11.0/24 [110/3] via 192.168.103.30, 19:45:41, GigabitEthernet0/0  
O IA 192.168.12.0/24  
[110/13] via 192.168.103.30, 19:45:41, GigabitEthernet0/0  
O IA 192.168.21.0/24  
[110/13] via 192.168.103.30, 19:45:41, GigabitEthernet0/0  
O IA 192.168.22.0/24 [110/3] via 192.168.103.30, 19:45:41, GigabitEthernet0/0  
O IA 192.168.31.0/24 [110/2] via 192.168.103.30, 19:45:41, GigabitEthernet0/0  
O IA 192.168.32.0/24 [110/2] via 192.168.103.30, 19:45:41, GigabitEthernet0/0  
O IA 192.168.41.0/24 [110/3] via 192.168.103.30, 19:41:28, GigabitEthernet0/0  
O IA 192.168.42.0/24 [110/3] via 192.168.103.30, 19:45:41, GigabitEthernet0/0  
O IA 192.168.104.0/24  
[110/3] via 192.168.103.30, 19:40:54, GigabitEthernet0/0  
O IA 192.168.134.0/24  
[110/2] via 192.168.103.30, 19:41:28, GigabitEthernet0/0
```

Output on the PE router should be similar to the following:

```
RP/0/RSP0/CPU0:PE3# show route ospf
```

```
Fri May 16 07:52:18.190 UTC
```

```
O 10.0.1.1/32 [110/2] via 192.168.31.1, 19:46:28, GigabitEthernet0/0/0/2  
O 10.0.2.1/32 [110/2] via 192.168.32.2, 19:46:23, GigabitEthernet0/0/0/3  
O 10.1.1.1/32 [110/3] via 192.168.31.1, 19:46:28, GigabitEthernet0/0/0/2  
O 10.2.1.1/32 [110/3] via 192.168.32.2, 19:46:23, GigabitEthernet0/0/0/3  
O 10.3.10.1/32 [110/2] via 192.168.103.31, 17:18:43, GigabitEthernet0/0/0/0  
O IA 10.4.1.1/32 [110/2] via 192.168.134.40, 19:41:59, GigabitEthernet0/0/0/1  
O IA 10.4.10.1/32 [110/3] via 192.168.134.40, 19:41:24, GigabitEthernet0/0/0/1  
O 192.168.1.0/24 [110/2] via 192.168.32.2, 19:46:23, GigabitEthernet0/0/0/3  
[110/2] via 192.168.31.1, 19:46:23, GigabitEthernet0/0/0/2  
O 192.168.2.0/24 [110/11] via 192.168.32.2, 19:46:23, GigabitEthernet0/0/0/3  
[110/11] via 192.168.31.1, 19:46:23, GigabitEthernet0/0/0/2  
O 192.168.11.0/24 [110/2] via 192.168.31.1, 19:46:28, GigabitEthernet0/0/0/2  
O 192.168.12.0/24 [110/12] via 192.168.31.1, 19:46:28, GigabitEthernet0/0/0/2  
O 192.168.21.0/24 [110/12] via 192.168.32.2, 19:46:23, GigabitEthernet0/0/0/3  
O 192.168.22.0/24 [110/2] via 192.168.32.2, 19:46:23, GigabitEthernet0/0/0/3  
O 192.168.41.0/24 [110/2] via 192.168.134.40, 19:41:59, GigabitEthernet0/0/0/1  
O 192.168.42.0/24 [110/2] via 192.168.32.2, 19:41:59, GigabitEthernet0/0/0/3  
[110/2] via 192.168.134.40, 19:41:59, GigabitEthernet0/0/0/1  
O IA 192.168.104.0/24 [110/2] via 192.168.134.40, 19:41:24, GigabitEthernet0/0/0/1
```

Step 4 Verify IPv4 connectivity between the CE routers in your team. A **ping** from the CE router in your pod to the CE router in the other pod (same team) should be successful.

```
CE3# ping 10.4.10.1 (IP address of the CE4 router Lo0 interface)
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.4.10.1, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

Step 5 On the CE router in your pod, verify that these IPv6 protocols are running: OSPFv3 neighbors and the OSPFv3 database. There should be one OSPFv3 neighbor. The OSPFv3 database should show LSAs from Area 1x or 1y, where x or y is your pod number. Outputs should be similar to the following:

```
CE3# show ipv6 protocols
IPv6 Routing Protocol is "connected"
IPv6 Routing Protocol is "ND"
IPv6 Routing Protocol is "ospf 1"
  Interfaces (Area 13):
    Loopback0
    GigabitEthernet0/0
  Redistribution:
    None
CE3# show ipv6 ospf neighbor

      OSPFv3 Router with ID (10.3.10.1) (Process ID 1)

Neighbor ID    Pri   State           Dead Time   Interface ID  Interface
10.3.1.1        1    FULL/BDR        00:00:35    9             GigabitEthernet0/0
CE3# show ipv6 ospf database

      OSPFv3 Router with ID (10.3.10.1) (Process ID 1)

      Router Link States (Area 13)

ADV Router    Age      Seq#          Fragment ID  Link count  Bits
10.3.1.1      166      0x80000023    0            1           B
10.3.10.1     1924     0x80000022    0            1           None

      Net Link States (Area 13)

ADV Router    Age      Seq#          Link ID      Rtr count
10.3.10.1     1924     0x80000021    4            2

      Inter Area Prefix Link States (Area 13)

ADV Router    Age      Seq#          Prefix
10.3.1.1      910      0x80000020    2001:DB8:10:3:1::1/128
10.3.1.1      910      0x80000020    2001:DB8:10:4:1::1/128
10.3.1.1      910      0x80000020    2001:DB8:10:4:10::1/128

      Link (Type-8) Link States (Area 13)

ADV Router    Age      Seq#          Link ID      Interface
10.3.1.1      419      0x80000020    9            Gi0/0
10.3.10.1     943      0x80000020    4            Gi0/0

      Intra Area Prefix Link States (Area 13)

ADV Router    Age      Seq#          Link ID      Ref-lstype  Ref-LSID
10.3.10.1     943      0x80000020    0            0x2001      0
```

- Step 6** On the PE router in your pod, verify that these IPv6 protocols are running: OSPFv3 neighbors and the OSPFv3 database. There should be four OSPFv3 neighbors. The OSPFv3 database should show LSAs from Areas 0, x, or y, and 1x or 1y, where x or y is your pod number. Outputs should be similar to the following:

```
RP/0/RSP0/CPU0:PE3# show protocols ipv6
```

```
Fri May 16 07:57:39.368 UTC
```

```
Routing Protocol OSPFv3 1
```

```
Router Id: 10.3.1.1
```

```
Distance: 110
```

```
Graceful Restart: Disabled
```

```
Redistribution:
```

```
None
```

```
Area 0
```

```
GigabitEthernet0/0/0/3
```

```
GigabitEthernet0/0/0/2
```

```
GigabitEthernet0/0/0/1
```

```
Area 3
```

```
Loopback0
```

```
Area 13
```

```
GigabitEthernet0/0/0/0
```

```
RP/0/RSP0/CPU0:PE3# show ospfv3 neighbor
```

```
Fri May 16 07:57:46.826 UTC
```

```
Neighbors for OSPFv3 1
```

Neighbor ID	Pri	State	Dead Time	Interface ID	Interface
10.3.10.1	1	FULL/DR	00:00:32	4	
GigabitEthernet0/0/0/0					
Neighbor is up for 17:24:14					
10.4.1.1	1	FULL/BDR	00:00:32	9	
GigabitEthernet0/0/0/1					
Neighbor is up for 17:34:08					

```
Total neighbor count: 2
```

```
RP/0/RSP0/CPU0:PE3# show ospfv3 database
```

```
Fri May 16 07:57:54.117 UTC
```

```
OSPFv3 Router with ID (10.3.1.1) (Process ID 1)
```

```
Router Link States (Area 0)
```

ADV Router	Age	Seq#	Fragment ID	Link count	Bits
10.3.1.1	995	0x80000021	0	1	B
10.4.1.1	713	0x80000021	0	1	B

```
Net Link States (Area 0)
```

ADV Router	Age	Seq#	Link ID	Rtr count
10.3.1.1	995	0x80000020	10	2

```
Inter Area Prefix Link States (Area 0)
```

ADV Router	Age	Seq#	Prefix
10.3.1.1	995	0x80000020	2001:db8:10:3:1::1/128
10.3.1.1	251	0x80000020	2001:db8:10:3:10::1/128
10.4.1.1	713	0x80000020	2001:db8:10:4:1::1/128
10.4.1.1	713	0x80000020	2001:db8:10:4:10::1/128

Link (Type-8) Link States (Area 0)

ADV Router	Age	Seq#	Link ID	Interface
10.3.1.1	995	0x80000020	10	Gi0/0/0/1
10.4.1.1	713	0x80000020	9	Gi0/0/0/1
10.3.1.1	995	0x80000020	11	Gi0/0/0/2
10.3.1.1	995	0x80000020	12	Gi0/0/0/3

Router Link States (Area 3)

ADV Router	Age	Seq#	Fragment ID	Link count	Bits
10.3.1.1	995	0x80000020	0	0	B

Inter Area Prefix Link States (Area 3)

ADV Router	Age	Seq#	Prefix
10.3.1.1	995	0x80000020	2001:db8:10:4:1::1/128
10.3.1.1	995	0x80000020	2001:db8:10:4:10::1/128
10.3.1.1	251	0x80000020	2001:db8:10:3:10::1/128

Intra Area Prefix Link States (Area 3)

ADV Router	Age	Seq#	Link ID	Ref-lstype	Ref-LSID
10.3.1.1	995	0x80000020	0	0x2001	0

Router Link States (Area 13)

ADV Router	Age	Seq#	Fragment ID	Link count	Bits
10.3.1.1	251	0x80000023	0	1	B
10.3.10.1	16	0x80000023	0	1	None

Net Link States (Area 13)

ADV Router	Age	Seq#	Link ID	Rtr count
10.3.10.1	16	0x80000022	4	2

Inter Area Prefix Link States (Area 13)

ADV Router	Age	Seq#	Prefix
10.3.1.1	995	0x80000020	2001:db8:10:3:1::1/128
10.3.1.1	995	0x80000020	2001:db8:10:4:1::1/128
10.3.1.1	995	0x80000020	2001:db8:10:4:10::1/128

Link (Type-8) Link States (Area 13)

ADV Router	Age	Seq#	Link ID	Interface
10.3.1.1	504	0x80000020	9	Gi0/0/0/0
10.3.10.1	1030	0x80000020	4	Gi0/0/0/0

Intra Area Prefix Link States (Area 13)

ADV Router	Age	Seq#	Link ID	Ref-lstype	Ref-LSID
10.3.10.1	1030	0x80000020	0	0x2001	0

Step 7 On the CE and PE routers, verify that OSPFv3 routes are present in the IPv6 routing table. Output on the CE router should be similar to the following:

```

CE3# show ipv6 route ospf
IPv6 Routing Table - default - 5 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
        B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
        H - NHRP, I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
        IS - ISIS summary, D - EIGRP, EX - EIGRP external, NM - NEMO
        ND - Neighbor Discovery, l - LISP
        O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
        ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
OI 2001:DB8:10:3:1::1/128 [110/1]
    via FE80::6E9C:EDFF:FE24:1C28, GigabitEthernet0/0
OI 2001:DB8:10:4:1::1/128 [110/2]
    via FE80::6E9C:EDFF:FE24:1C28, GigabitEthernet0/0
OI 2001:DB8:10:4:10::1/128 [110/3]
    via FE80::6E9C:EDFF:FE24:1C28, GigabitEthernet0/0

```

Output on the PE router should be similar to the following:

```

RP/0/RSP0/CPU0:PE3# show route ipv6 ospf
Fri May 16 08:00:11.084 UTC

O    2001:db8:10:3:10::1/128
    [110/1] via fe80::2e54:2dff:fe9b:d3e0, 17:26:37, GigabitEthernet0/0/0/0
O IA 2001:db8:10:4:1::1/128
    [110/1] via fe80::462b:3ff:fee7:9381, 17:36:27, GigabitEthernet0/0/0/1
O IA 2001:db8:10:4:10::1/128
    [110/2] via fe80::462b:3ff:fee7:9381, 17:36:27, GigabitEthernet0/0/0/1

```

Step 8 Verify IPv6 connectivity between the CE routers in your team. A **ping** from the CE router in your pod to the CE router in the other pod (same team) should be successful.

```

CE3# ping 2001:db8:10:4:10::1 {IPv6 address of the CE2 router Lo0 interface}
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:10:4:10::1, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 0/2/8 ms

```

Task 2: Influencing OSPF DR and BDR Election

In this task, you will influence OSPF DR and BDR election on the LAN. On the LAN segment between the CE and PE routers, only the PE router will become the DR, and the CE router will neither become the DR nor BDR. On the LAN segments between the PE router and the core routers, you will influence DR and BDR election by changing the OSPF priority.

Note Outputs in this task are taken from team 2 with pods 3 and 4.

Activity Procedure

Complete these steps:

- Step 1** On the CE router Gi0/0 interface, change the OSPF priority to the value that enables the PE router Gi0/0/0/0 interface to be the DR.
- Step 2** On the PE router Gi0/0/0/2 and Gi0/0/0/3 interfaces, change the OSPF priority to the value that is one higher than the default.

Verification

You have completed this task when you attain these results:

- Step 1** On the PE router in your pod, verify the state of the Gi0/0/0/0 interface. The state should be the DR.

```
RP/0/RSP0/CPU0:PE3# show ospf interface GigabitEthernet 0/0/0/0 | include State
Fri May 16 08:14:38.951 UTC
Transmit Delay is 1 sec, State DR, Priority 1, MTU 1500, MaxPktSz 1500
```

- Step 2** On the PE router in your pod, verify the priority that the Gi0/0/0/2 and Gi0/0/0/3 interfaces have. The priority should be 2.

```
RP/0/RSP0/CPU0:PE3# show ospf interface GigabitEthernet 0/0/0/2 | include State
Fri May 16 08:15:27.583 UTC
Transmit Delay is 1 sec, State BDR, Priority 2, MTU 1500, MaxPktSz 1500
```

```
RP/0/RSP0/CPU0:PE3# show ospf interface GigabitEthernet 0/0/0/3 | include State
Fri May 16 08:15:46.953 UTC
Transmit Delay is 1 sec, State BDR, Priority 2, MTU 1500, MaxPktSz 1500
```

If the state on the Gi0/0/0/2 and Gi0/0/0/3 interfaces is *not* the DR, it should change when the current DR will go down.

Task 3: Influencing OSPF Route Selection by Changing the OSPF Link Cost

In this task, you will change the OSPF link cost to influence the OSPF route selection. The OSPF that is configured on the PE router in your pod selects the best route to the neighbor pod (same team) via a link that directly connects both pods. You will change the OSPF cost on the link to the value that will make OSPF select the path via core devices.

Note	Outputs in this task are taken from team 2 with pods 3 and 4.
-------------	---

Activity Procedure

Complete these steps:

- Step 1** On the PE router, verify the route that is selected for prefixes from the neighbor pod. The OSPF selects a direct link between pods and does not use backbone routers.

```
RP/0/RSP0/CPU0:PE3# show route 10.4.1.1
Fri May 16 08:35:05.187 UTC

Routing entry for 10.4.1.1/32
  Known via "ospf 1", distance 110, metric 2, type inter area
  Installed May 15 12:10:18.822 for 20:24:46
  Routing Descriptor Blocks
    192.168.134.40, from 4.4.4.4, via GigabitEthernet0/0/0/1
    Route metric is 2
  No advertising protos.
```

```
RP/0/RSP0/CPU0:PE3# show route 10.4.10.1
Fri May 16 08:36:06.125 UTC

Routing entry for 10.4.10.1/32
  Known via "ospf 1", distance 110, metric 3, type inter area
  Installed May 15 12:10:53.802 for 20:25:12
  Routing Descriptor Blocks
    192.168.134.40, from 4.4.4.4, via GigabitEthernet0/0/0/1
    Route metric is 3
  No advertising protos.
```

- Step 2** On the PE (Cisco IOS XR) router, change the OSPF link cost on the Gi0/0/0/3 (to P2) to 3. On the PE (Cisco IOS XE) router, change the OSPF link cost on the Gi0/0/2 (to P1) to 3.
- Step 3** On the PE routers in your team, change the OSPF link cost on the direct link between pods to 10.

Verification

You have completed this task when you attain these results:

- Step 1** The PE router in your pod selects the best path to the prefixes in the neighbor pod via core routers. The routing table on the PE (Cisco IOS XR) router shows the best path via the P1 core router:

```
RP/0/RSP0/CPU0:PE3# show route 10.4.1.1
Fri May 16 10:24:02.051 UTC

Routing entry for 10.4.1.1/32
  Known via "ospf 1", distance 110, metric 4, type inter area
  Installed May 16 10:22:33.188 for 00:01:29
  Routing Descriptor Blocks
    192.168.31.1, from 4.4.4.4, via GigabitEthernet0/0/0/2
    Route metric is 4
  No advertising protos.
```

```
RP/0/RSP0/CPU0:PE3# show route 10.4.10.1
Fri May 16 10:24:25.051 UTC

Routing entry for 10.4.10.1/32
  Known via "ospf 1", distance 110, metric 5, type inter area
  Installed May 16 10:22:33.188 for 00:01:52
  Routing Descriptor Blocks:
    192.168.31.1, from 4.4.4.4, via GigabitEthernet0/0/0/2
      Route metric is 5
  No advertising protos.
```

The routing table on the PE (Cisco IOS XE) router shows the best path via the P2 core router:

```
PE4# show ip route 10.3.1.1
Routing entry for 10.3.1.1/32
  Known via "ospf 1", distance 110, metric 4, type inter area
  Last update from 192.168.42.2 on GigabitEthernet0/0/3, 00:02:07 ago
  Routing Descriptor Blocks:
    * 192.168.42.2, from 3.3.3.3, 00:02:07 ago, via GigabitEthernet0/0/3
      Route metric is 4, traffic share count is 1
```

```
PE4# show ip route 10.3.10.1
Routing entry for 10.3.10.1/32
  Known via "ospf 1", distance 110, metric 5, type inter area
  Last update from 192.168.42.2 on GigabitEthernet0/0/3, 00:02:11 ago
  Routing Descriptor Blocks:
    * 192.168.42.2, from 3.3.3.3, 00:02:11 ago, via GigabitEthernet0/0/3
      Route metric is 5, traffic share count is 1
```

Step 2 The trace between the PE routers in your team shows the core routers in the path. Output on the PE (Cisco IOS XR) router should be similar to the following:

```
RP/0/RSP0/CPU0:PE3# traceroute 10.4.1.1
Fri May 16 10:26:47.485 UTC

Type escape sequence to abort.
Tracing the route to 10.4.1.1

 1 192.168.31.1 1 msec 0 msec 0 msec
 2 192.168.1.2 [MPLS: Label 16012 Exp 0] 0 msec 0 msec 0 msec
 3 192.168.42.40 0 msec * 0 msec
```

Output on the PE (Cisco IOS XE) router should be similar to the following:

```
PE4# traceroute 10.3.1.1
Type escape sequence to abort.
Tracing the route to 10.3.1.1
VRF info: (vrf in name/id, vrf out name/id)
 1 192.168.42.2 1 msec 1 msec 1 msec
 2 192.168.1.1 [MPLS: Label 16004 Exp 0] 1 msec 1 msec 1 msec
 3 192.168.31.30 52 msec 2 msec *
```

Task 4: Configuring OSPF Authentication

In this task, you will configure OSPF authentication between the routers in your pod.

Note Outputs in this task are taken from team 2 with pods 3 and 4.

Activity Procedure

Complete these steps:

Step 1 On the CE and PE routers, enable OSPF MD5 authentication. Use key number “1” and key string “cisco.”

Verification

You have completed this task when you attain these results:

Step 1 On the PE router in your pod, verify that the OSPF neighbor to the CE router is up. Output from the PE (Cisco IOS XR) router should be similar to the following:

```
RP/0/RSP0/CPU0:PE3# show ospf neighbor 30.30.30.30 | include State
Fri May 16 10:49:12.001 UTC
Neighbor priority is 0, State is FULL, 6 state changes
```

Output from the PE (Cisco IOS XE) router should be similar to the following:

```
PE4# show ip ospf neighbor 40.40.40.40 | include State
Neighbor priority is 0, State is FULL, 6 state changes
```

Step 2 On the PE router in your pod, verify that OSPF MD5 authentication is used on the interface toward the CE router. Output from the PE (Cisco IOS XR) router should be similar to the following:

```
RP/0/RSP0/CPU0:PE3# show ospf interface GigabitEthernet 0/0/0/0
Fri May 16 10:51:04.527 UTC

GigabitEthernet0/0/0/0 is up, line protocol is up
  Internet Address 192.168.103.30/24, Area 13
  Process ID 1, Router ID 3.3.3.3, Network Type BROADCAST, Cost: 1
  Transmit Delay is 1 sec, State DR, Priority 1, MTU 1500, MaxPktSz 1500
  Designated Router (ID) 3.3.3.3, Interface address 192.168.103.30
  No backup designated router on this network
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:07
  Index 1/5, flood queue length 0
  Next 0(0)/0(0)
  Last flood scan length is 1, maximum is 21
  Last flood scan time is 0 msec, maximum is 0 msec
  LS Ack List: current length 0, high water mark 10
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 30.30.30.30
  Suppress hello for 0 neighbor(s)
  Message digest authentication enabled
  Youngest key id is 1
  Multi-area interface Count is 0
```

Output from the PE (Cisco IOS XE) router should be similar to the following:

```
PE4# show ip ospf interface GigabitEthernet 0/0/0
GigabitEthernet0/0/0 is up, line protocol is up
  Internet Address 192.168.104.40/24, Area 14, Attached via Network Statement
  Process ID 1, Router ID 4.4.4.4, Network Type BROADCAST, Cost: 1
  Topology-MTID      Cost      Disabled      Shutdown      Topology Name
    0                1         no            no            Base
  Transmit Delay is 1 sec, State DR, Priority 1
  Designated Router (ID) 4.4.4.4, Interface address 192.168.104.40
  No backup designated router on this network
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    oob-resync timeout 40
    Hello due in 00:00:06
  Supports Link-local Signaling (LLS)
  Cisco NSF helper support enabled
  IETF NSF helper support enabled
  Index 1/4, flood queue length 0
  Next 0x0(0)/0x0(0)
  Last flood scan length is 1, maximum is 20
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 40.40.40.40
  Suppress hello for 0 neighbor(s)
  Message digest authentication enabled
  Youngest key id is 1
```

Task 5: Configuring OSPF Virtual Links

In this task, you will configure an OSPF virtual link to connect the nonbackbone area to the backbone area. For the purpose of this task, you will introduce another network on the CE router and place this network into the OSPF nonbackbone area. To reach the newly introduced network, you will enable the OSPF virtual link between the CE and PE routers.

Note Outputs in this task are taken from team 2 with pods 3 and 4.

Activity Procedure

Complete these steps:

- Step 1** On the CE router, enable the Loopback1x or Loopback1y interface and assign IP address 1x.0.0.1/24 or 1y.0.0.1/24, where x or y is your pod number.
- Step 2** On the CE router, include Loopback1x or Loopback1y into OSPF Area x00 or y00.
- Step 3** Configure the OSPF virtual link between the CE and PE routers.

Verification

You have completed this task when you attain these results:

- Step 1** On the CE router in your pod, verify that the OSPF virtual link is running and OSPF adjacency is up on the virtual link.

```
CE3# show ip ospf virtual-links
Virtual Link OSPF VL0 to router 3.3.3.3 is up
  Run as demand circuit
  DoNotAge LSA allowed.
  Transit area 13, via interface GigabitEthernet0/0
Topology-MTID   Cost      Disabled   Shutdown   Topology Name
   0             1         no         no         Base
Transmit Delay is 1 sec, State POINT TO POINT,
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
Hello due in 00:00:00
Adjacency State FULL (Hello suppressed)
Index 1/2, retransmission queue length 0, number of retransmission 1
First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)
Last retransmission scan length is 1, maximum is 1
Last retransmission scan time is 0 msec, maximum is 0 msec
```

```
CE3# show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
3.3.3.3	0	FULL/ -	-	192.168.103.30	OSPF VL0
3.3.3.3	1	FULL/DR	00:00:39	192.168.103.30	GigabitEthernet0/0

Step 2 On the CE router in your pod, verify that the OSPF database has LSAs from Area 0.

CE3# show ip ospf database

OSPF Router with ID (30.30.30.30) (Process ID 1)

Router Link States (Area 0)

Link ID	ADV Router	Age		Seq#	Checksum	Link count
0.0.0.1	0.0.0.1	652	{DNA}	0x800002EB	0x00B52A	7
0.0.0.2	0.0.0.2	110	{DNA}	0x800002EE	0x00BE13	7
1.1.1.1	1.1.1.1	619	{DNA}	0x800002D7	0x00E640	3
2.2.2.2	2.2.2.2	566	{DNA}	0x80000281	0x008992	3
3.3.3.3	3.3.3.3	1	{DNA}	0x80000036	0x006452	4
4.4.4.4	4.4.4.4	2	{DNA}	0x8000002F	0x00D764	4
30.30.30.30	30.30.30.30	270		0x80000002	0x00A49D	1
40.40.40.40	40.40.40.40	3	{DNA}	0x80000002	0x00657D	1

Net Link States (Area 0)

Link ID	ADV Router	Age		Seq#	Checksum
192.168.1.2	0.0.0.2	110	{DNA}	0x800002CC	0x004BAF
192.168.2.2	0.0.0.2	110	{DNA}	0x800002CC	0x0040B9
192.168.11.10	1.1.1.1	619	{DNA}	0x800002CD	0x008E55
192.168.12.2	0.0.0.2	110	{DNA}	0x800002CA	0x00F5F8
192.168.21.20	2.2.2.2	566	{DNA}	0x8000027F	0x0050C5
192.168.22.20	2.2.2.2	566	{DNA}	0x800002CD	0x00B60F
192.168.31.1	0.0.0.1	129	{DNA}	0x80000034	0x00C1AD
192.168.32.2	0.0.0.2	610	{DNA}	0x80000034	0x00B0BA
192.168.41.1	0.0.0.1	1906	{DNA}	0x8000002F	0x008FD6
192.168.42.2	0.0.0.2	110	{DNA}	0x80000030	0x007CE4
192.168.134.40	4.4.4.4	740	{DNA}	0x80000002	0x0039BB

Summary Net Link States (Area 0)

Link ID	ADV Router	Age		Seq#	Checksum
10.3.1.1	3.3.3.3	572	{DNA}	0x8000002B	0x0029CD
10.3.10.1	3.3.3.3	944	{DNA}	0x80000001	0x0024F2
10.3.10.1	30.30.30.30	322		0x80000002	0x00EABF
10.4.1.1	4.4.4.4	8	{DNA}	0x8000002A	0x0001F1
10.4.10.1	4.4.4.4	915	{DNA}	0x80000001	0x00F918
10.4.10.1	40.40.40.40	91	{DNA}	0x80000002	0x00B1CF
13.0.0.1	30.30.30.30	322		0x80000002	0x00565E
14.0.0.1	40.40.40.40	91	{DNA}	0x80000002	0x001C6F
192.168.103.0	3.3.3.3	572	{DNA}	0x80000031	0x00B579
192.168.103.0	30.30.30.30	322		0x80000002	0x00E60B
192.168.104.0	4.4.4.4	1499	{DNA}	0x80000029	0x009C95
192.168.104.0	40.40.40.40	91	{DNA}	0x80000002	0x00AE1A

Type-10 Opaque Link Area Link States (Area 0)

Link ID	ADV Router	Age		Seq#	Checksum	Opaque ID
1.0.0.0	0.0.0.1	652	{DNA}	0x800002CD	0x00084E	0
1.0.0.0	0.0.0.2	855	{DNA}	0x800002AC	0x004D28	0
1.0.0.9	0.0.0.2	855	{DNA}	0x800002AD	0x004C02	9
1.0.0.10	0.0.0.2	109	{DNA}	0x800002AF	0x000146	10
1.0.0.11	0.0.0.2	109	{DNA}	0x800002AF	0x007DB2	11
1.0.0.12	0.0.0.1	1906	{DNA}	0x800002CB	0x009F14	12
1.0.0.12	0.0.0.2	109	{DNA}	0x8000002F	0x006623	12
1.0.0.13	0.0.0.1	1404	{DNA}	0x800002CD	0x005458	13
1.0.0.13	0.0.0.2	855	{DNA}	0x8000027D	0x00300E	13
1.0.0.14	0.0.0.1	652	{DNA}	0x800002CD	0x002E65	14
1.0.0.14	0.0.0.2	855	{DNA}	0x8000002E	0x00294F	14
1.0.0.15	0.0.0.1	1159	{DNA}	0x8000002E	0x00C8C7	15
1.0.0.16	0.0.0.1	1404	{DNA}	0x8000027D	0x00E852	16
1.0.0.17	0.0.0.1	1159	{DNA}	0x8000002F	0x00C9AB	17

Router Link States (Area 13)

Link ID	ADV Router	Age	Seq#	Checksum	Link count
3.3.3.3	3.3.3.3	843	0x80000005	0x0072C5	1
30.30.30.30	30.30.30.30	270	0x80000035	0x00D62A	2

Net Link States (Area 13)

Link ID	ADV Router	Age	Seq#	Checksum
192.168.103.30	3.3.3.3	1216	0x80000002	0x003B7E

Summary Net Link States (Area 13)

Link ID	ADV Router	Age	Seq#	Checksum
10.0.1.1	3.3.3.3	843	0x80000002	0x00A978
10.0.2.1	3.3.3.3	843	0x80000002	0x00A877
10.1.1.1	3.3.3.3	843	0x80000002	0x00A778
10.2.1.1	3.3.3.3	843	0x80000002	0x00A578
10.3.1.1	3.3.3.3	843	0x80000002	0x007BA4
10.4.1.1	3.3.3.3	843	0x80000003	0x008B8F
10.4.10.1	3.3.3.3	1184	0x80000001	0x0036DC
13.0.0.1	3.3.3.3	271	0x80000001	0x008F91
13.0.0.1	30.30.30.30	332	0x80000001	0x00585D
14.0.0.1	3.3.3.3	224	0x80000001	0x00A07C
192.168.1.0	3.3.3.3	843	0x80000002	0x00843F
192.168.2.0	3.3.3.3	843	0x80000002	0x00833E
192.168.11.0	3.3.3.3	843	0x80000002	0x0016A3
192.168.12.0	3.3.3.3	843	0x80000002	0x00298C
192.168.21.0	3.3.3.3	843	0x80000002	0x00BBF1
192.168.22.0	3.3.3.3	843	0x80000002	0x00A607
192.168.31.0	3.3.3.3	843	0x80000002	0x002F77
192.168.32.0	3.3.3.3	843	0x80000004	0x00346D
192.168.41.0	3.3.3.3	843	0x80000002	0x00DEBA
192.168.42.0	3.3.3.3	843	0x80000003	0x00C7D0
192.168.104.0	3.3.3.3	843	0x80000003	0x002534
192.168.134.0	3.3.3.3	843	0x80000006	0x001022

Router Link States (Area 300)

Link ID	ADV Router	Age	Seq#	Checksum	Link count
30.30.30.30	30.30.30.30	331	0x80000002	0x0070C0	1

Summary Net Link States (Area 300)

Link ID	ADV Router	Age	Seq#	Checksum
10.0.1.1	30.30.30.30	265	0x80000001	0x00882D
10.0.2.1	30.30.30.30	265	0x80000001	0x00872C
10.1.1.1	30.30.30.30	265	0x80000001	0x00862D
10.2.1.1	30.30.30.30	265	0x80000001	0x00842D
10.3.1.1	30.30.30.30	265	0x80000001	0x005A59
10.3.10.1	30.30.30.30	332	0x80000001	0x00ECBE
10.4.1.1	30.30.30.30	265	0x80000001	0x006C43
10.4.10.1	30.30.30.30	265	0x80000001	0x001392
14.0.0.1	30.30.30.30	208	0x80000001	0x007D32
192.168.1.0	30.30.30.30	265	0x80000001	0x0063F3
192.168.2.0	30.30.30.30	265	0x80000001	0x0062F2
192.168.11.0	30.30.30.30	265	0x80000001	0x00F458
192.168.12.0	30.30.30.30	265	0x80000001	0x000841
192.168.21.0	30.30.30.30	265	0x80000001	0x009AA6
192.168.22.0	30.30.30.30	265	0x80000001	0x0085BB
192.168.31.0	30.30.30.30	265	0x80000001	0x000E2C
192.168.32.0	30.30.30.30	265	0x80000001	0x001720
192.168.41.0	30.30.30.30	265	0x80000001	0x00BD6F
192.168.42.0	30.30.30.30	265	0x80000001	0x00A884
192.168.103.0	30.30.30.30	332	0x80000001	0x00E80A
192.168.104.0	30.30.30.30	265	0x80000001	0x0006E7
192.168.134.0	30.30.30.30	265	0x80000001	0x00F6D2

Step 3 Verify IP connectivity from the PE router in your pod to the newly enabled Loopback IP address on the CE router in your pod.

```
RP/0/RSP0/CPU0:PE3# ping 13.0.0.1
Fri May 16 11:08:47.731 UTC
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 13.0.0.1, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

Hardware Lab 2: Implement OSPF Special Area Types

Activity Objective

In this lab activity, you will configure OSPF summarization. You will configure the nonbackbone OSPF area to be a stub and totally stubby area. You will also adjust the cost of the injected default route on an ABR router.


Note	Students from two different pods are working as 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 running Cisco IOS XR Software, and the second pod in the same team will work on the PE router running Cisco IOS XE Software. Students in the same team should coordinate their lab activity.
-------------	---

In the lab activity, you will work on different Cisco routers 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:

- Summarize the subnets in the nonbackbone OSPF area
- Configure the nonbackbone OSPF area to be a stub area
- Configure the nonbackbone OSPF area to be a totally stubby area
- Adjust the cost of the injected default route on an ABR router
- Configure the nonbackbone OSPF area to be a totally NSSA

Visual Objective

The figure illustrates what you will accomplish in this activity.



OSPF Areas

This section includes a table with OSPF areas that are used in the lab. Replace the x or y with your pod number to get the OSPF areas that are used in the lab.

Link	OSPF area
P1 – P2	0
P1 and P2 Loopback0	
PEx – P1, PEy – P1	
PEx – P2, PEy – P2	
PEx – PEy	
PEx Loopback0, PEy Loopback0	x, y
CEx – PEx, CEy – PEy	1x, 1y
CEx Loopback0, CEy Loopback0	
CEx Loopback1x (1x.0.0.1/24), CEy Loopback1y (1y.0.0.1/24)	x00, y00

Required Resources

These are the resources and equipment that 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] area <i>area</i> virtual-link <i>IP-address</i>	(Router) Enables or disables the OSPF virtual link.
[no] network <i>prefix wildcard_mask area area</i>	(Router) Enables or disables OSPF routing on the network.
area <i>area</i> default-cost <i>cost</i>	(Router) Changes the cost of the default route that is sent by the ABR into the stub area.
area <i>area</i> range <i>prefix subnet-mask</i>	(Router) Summarizes prefixes from the nonbackbone OSPF area to the backbone OSPF area.
area <i>area</i> stub nssa [no-summary]	(Router) Changes the nonbackbone OSPF area into the stub area or NSSA. Use the no-summary keyword to create the totally stubby area or totally NSSA.
configure terminal	Enters configuration mode.
match interface <i>interface</i>	(Route-map) Creates a statement to match the interface in the route map.
ping <i>dest_IP</i> source [<i>source_IP</i>]	Verifies connectivity between the source IP and the destination IP.
redistribute connected [subnets] [route-map <i>route-map</i>]	(Router) Redistributes connected routes into OSPF. Use the subnets keyword to include subnets into redistribution. Use the route-map keyword to redistribute with conditions.
route-map <i>name</i> permit deny	Creates the route-map sentence and enters route-map configuration mode.
router ospf <i>process-id</i>	Enables the OSPF process on the router.
show ip ospf database	Displays the content of the OSPF database.
show ip route	Displays the current routes in the routing table.
summary-address <i>prefix subnet-mask</i>	(Router) Summarizes external OSPF prefixes into the OSPF domain.

Cisco IOS XR Commands

Command	Description
[no] virtual-link <i>IP-address</i>	(Router) Enables or disables the OSPF virtual link.

Command	Description
area <i>area</i>	(Router) Creates the OSPF area on the router.
commit	Commits changes to the running configuration.
configure terminal	Enters configuration mode.
default-cost <i>cost</i>	(Router) Changes the cost of the default route that is sent by the ABR into the stub area.
ping <i>dest_IP</i> <i>source</i> [<i>source_IP</i>]	Verifies connectivity between the source IP and the destination IP.
range <i>prefix/subnet-mask</i>	(Router) Summarizes prefixes from the nonbackbone OSPF area to the backbone OSPF area.
router ospf <i>process-ID</i>	Creates an OSPF process.
show ospf database	Displays the content of the OSPF database.
show route	Displays the current routes in the routing table.
stub nssa [no-summary]	(Router) Changes the nonbackbone OSPF area into the stub area or NSSA. Use the no-summary keyword to create the totally stubby area or totally NSSA.

Task 1: Enabling OSPF Summarization

In this task, you will enable OSPF summarization.

Outputs in this task are taken from team 2 with pods 3 and 4.

Activity Procedure

Complete these steps:

- Step 1** On the CE and PE routers, remove the OSPF virtual link that is configured in the previous lab activity and, on the CE router, remove Loopback1x or Loopback1y from the OSPF process.
- Step 2** On the CE router, redistribute Loopback1x or Loopback1y into the OSPF process. You will see that the external route appears in the PE router routing table.
- On the PE router, you will see two OSPF routes from the CE router: Loopback0 and Loopback1x or Loopback1y, where x or y is your pod number. The CE router Loopback0 network is shown as the intra-area OSPF route, and the Loopback1x or Loopback1y network is shown as the external OSPF route.

```
RP/0/RSP0/CPU0:PE3# show route ospf
```

```
Fri May 16 12:14:27.953 UTC
```

```
O 10.0.1.1/32 [110/2] via 192.168.31.1, 01:54:01, GigabitEthernet0/0/0/2
O 10.0.2.1/32 [110/3] via 192.168.31.1, 01:51:54, GigabitEthernet0/0/0/2
O 10.1.1.1/32 [110/3] via 192.168.31.1, 01:54:01, GigabitEthernet0/0/0/2
O 10.2.1.1/32 [110/4] via 192.168.31.1, 01:51:54, GigabitEthernet0/0/0/2
O 10.3.10.1/32 [110/2] via 192.168.103.31, 01:27:13, GigabitEthernet0/0/0/0
O IA 10.4.1.1/32 [110/4] via 192.168.31.1, 01:51:54, GigabitEthernet0/0/0/2
O IA 10.4.10.1/32 [110/5] via 192.168.31.1, 01:26:42, GigabitEthernet0/0/0/2
O E2 13.0.0.0/24 [110/20] via 192.168.103.31, 00:01:31, GigabitEthernet0/0/0/0
O E2 14.0.0.0/24 [110/20] via 192.168.31.1, 00:01:11, GigabitEthernet0/0/0/2
O 192.168.1.0/24 [110/2] via 192.168.31.1, 01:54:01, GigabitEthernet0/0/0/2
O 192.168.2.0/24 [110/3] via 192.168.31.1, 01:54:01, GigabitEthernet0/0/0/2
O 192.168.11.0/24 [110/2] via 192.168.31.1, 01:54:01, GigabitEthernet0/0/0/2
O 192.168.12.0/24 [110/5] via 192.168.31.1, 01:51:54, GigabitEthernet0/0/0/2
O 192.168.21.0/24 [110/4] via 192.168.31.1, 01:54:01, GigabitEthernet0/0/0/2
O 192.168.22.0/24 [110/3] via 192.168.31.1, 01:51:54, GigabitEthernet0/0/0/2
O 192.168.41.0/24 [110/4] via 192.168.31.1, 01:51:54, GigabitEthernet0/0/0/2
O 192.168.42.0/24 [110/3] via 192.168.31.1, 01:51:54, GigabitEthernet0/0/0/2
O IA 192.168.104.0/24 [110/4] via 192.168.31.1, 01:51:54, GigabitEthernet0/0/0/2
```

Step 3 On the PE router, summarize the intra-area OSPF (CE Loopback0 interface network) route to subnet mask 24.

Step 4 On the CE router, summarize the external OSPF (CE Loopback1x or Loopback1y interface network) route to subnet mask 8.

Verification

You have completed this task when you attain these results:

Step 1 On the PE router in your pod, verify the routing table. You will still see the intra-area OSPF route from the CE router with subnet mask 32 because the PE router is the ABR in Area 1x or 1y (where x or y is your pod number). Additionally, you will see a route pointing to the Null0 interface. The external OSPF route is shown with subnet mask 8. Outputs should be similar to the following:

```
RP/0/RSP0/CPU0:PE3# show route ospf
```

```
Fri May 16 12:21:54.417 UTC
```

```
O 10.0.1.1/32 [110/2] via 192.168.31.1, 02:01:27, GigabitEthernet0/0/0/2
O 10.0.2.1/32 [110/3] via 192.168.31.1, 01:59:21, GigabitEthernet0/0/0/2
O 10.1.1.1/32 [110/3] via 192.168.31.1, 02:01:27, GigabitEthernet0/0/0/2
O 10.2.1.1/32 [110/4] via 192.168.31.1, 01:59:21, GigabitEthernet0/0/0/2
O IA 10.3.10.0/24 [254/0] via 0.0.0.0, 00:04:18, Null0
O 10.3.10.1/32 [110/2] via 192.168.103.31, 01:34:40, GigabitEthernet0/0/0/0
O IA 10.4.1.1/32 [110/4] via 192.168.31.1, 01:59:21, GigabitEthernet0/0/0/2
O IA 10.4.10.0/24 [110/5] via 192.168.31.1, 00:03:27, GigabitEthernet0/0/0/2
O E2 13.0.0.0/8 [110/20] via 192.168.103.31, 00:01:56, GigabitEthernet0/0/0/0
O E2 14.0.0.0/8 [110/20] via 192.168.31.1, 00:01:43, GigabitEthernet0/0/0/2
O 192.168.1.0/24 [110/2] via 192.168.31.1, 02:01:27, GigabitEthernet0/0/0/2
O 192.168.2.0/24 [110/3] via 192.168.31.1, 02:01:27, GigabitEthernet0/0/0/2
O 192.168.11.0/24 [110/2] via 192.168.31.1, 02:01:27, GigabitEthernet0/0/0/2
O 192.168.12.0/24 [110/5] via 192.168.31.1, 01:59:21, GigabitEthernet0/0/0/2
O 192.168.21.0/24 [110/4] via 192.168.31.1, 02:01:27, GigabitEthernet0/0/0/2
O 192.168.22.0/24 [110/3] via 192.168.31.1, 01:59:21, GigabitEthernet0/0/0/2
O 192.168.41.0/24 [110/4] via 192.168.31.1, 01:59:21, GigabitEthernet0/0/0/2
O 192.168.42.0/24 [110/3] via 192.168.31.1, 01:59:21, GigabitEthernet0/0/0/2
O IA 192.168.104.0/24 [110/4] via 192.168.31.1, 01:59:21, GigabitEthernet0/0/0/2
```

Step 2 On the CE router in your pod, verify the routing table. You should see the CE router Loopback interface networks from the neighbor pod (same team). The subnet mask of Loopback0 is 24, and the subnet mask of Loopback 1y or Loopback 1x is 8 (where y or x is the neighbor pod number). Outputs should be similar to the following:

```
CE3# show ip route ospf
```

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

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

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

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

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

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

```
o - ODR, P - periodic downloaded static route, H - NHRP, I - LISP
```

```
+ - replicated route, % - next hop override
```

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

```

10.0.0.0/8 is variably subnetted, 8 subnets, 2 masks
O IA 10.0.1.1/32 [110/3] via 192.168.103.30, 00:17:17, GigabitEthernet0/0
O IA 10.0.2.1/32 [110/4] via 192.168.103.30, 00:17:17, GigabitEthernet0/0
O IA 10.1.1.1/32 [110/4] via 192.168.103.30, 00:17:17, GigabitEthernet0/0
O IA 10.2.1.1/32 [110/5] via 192.168.103.30, 00:17:17, GigabitEthernet0/0
O IA 10.3.1.1/32 [110/2] via 192.168.103.30, 00:17:17, GigabitEthernet0/0
O IA 10.4.1.1/32 [110/5] via 192.168.103.30, 00:17:17, GigabitEthernet0/0
O IA 10.4.10.0/24 [110/6] via 192.168.103.30, 00:06:08, GigabitEthernet0/0
13.0.0.0/8 is variably subnetted, 3 subnets, 3 masks
O 13.0.0.0/8 is a summary, 00:04:41, Null0
O E2 14.0.0.0/8 [110/20] via 192.168.103.30, 00:04:24, GigabitEthernet0/0
O IA 192.168.1.0/24 [110/3] via 192.168.103.30, 00:17:17, GigabitEthernet0/0
O IA 192.168.2.0/24 [110/4] via 192.168.103.30, 00:17:17, GigabitEthernet0/0
O IA 192.168.11.0/24 [110/3] via 192.168.103.30, 00:17:17, GigabitEthernet0/0
O IA 192.168.12.0/24 [110/6] via 192.168.103.30, 00:17:17, GigabitEthernet0/0
O IA 192.168.21.0/24 [110/5] via 192.168.103.30, 00:17:17, GigabitEthernet0/0
O IA 192.168.22.0/24 [110/4] via 192.168.103.30, 00:17:17, GigabitEthernet0/0
O IA 192.168.31.0/24 [110/2] via 192.168.103.30, 00:17:17, GigabitEthernet0/0
O IA 192.168.32.0/24 [110/4] via 192.168.103.30, 00:17:17, GigabitEthernet0/0
O IA 192.168.41.0/24 [110/5] via 192.168.103.30, 00:17:17, GigabitEthernet0/0
O IA 192.168.42.0/24 [110/4] via 192.168.103.30, 00:17:17, GigabitEthernet0/0
O IA 192.168.104.0/24
[110/5] via 192.168.103.30, 00:17:17, GigabitEthernet0/0
O IA 192.168.134.0/24
[110/11] via 192.168.103.30, 00:17:17, GigabitEthernet0/0

```

Task 2: Configuring the OSPF Stub Area

In this task, you will configure the nonbackbone OSPF area to be a stub and totally stubby area. You will also adjust the cost of the injected default route on an ABR router. To get full IP reachability to the external routes, you will enable NSSA.

Note Outputs in this task are taken from team 2 with pods 3 and 4.

Activity Procedure

Complete these steps:

Step 1 On the CE router, examine the IP routing table, and you should see the intra-area, interarea, and external OSPF routes. Output should be similar to the following:

```
CE3# show ip route
```

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

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

```
10.0.0.0/8 is variably subnetted, 8 subnets, 2 masks
O IA  10.0.1.1/32 [110/3] via 192.168.103.30, 00:21:39, GigabitEthernet0/0
O IA  10.0.2.1/32 [110/4] via 192.168.103.30, 00:21:39, GigabitEthernet0/0
O IA  10.1.1.1/32 [110/4] via 192.168.103.30, 00:21:39, GigabitEthernet0/0
O IA  10.2.1.1/32 [110/5] via 192.168.103.30, 00:21:39, GigabitEthernet0/0
O IA  10.3.1.1/32 [110/2] via 192.168.103.30, 00:21:39, GigabitEthernet0/0
C     10.3.10.1/32 is directly connected, Loopback0
O IA  10.4.1.1/32 [110/5] via 192.168.103.30, 00:21:39, GigabitEthernet0/0
O IA  10.4.10.0/24 [110/6] via 192.168.103.30, 00:10:30, GigabitEthernet0/0
13.0.0.0/8 is variably subnetted, 3 subnets, 3 masks
O     13.0.0.0/8 is a summary, 00:09:03, Null0
C     13.0.0.0/24 is directly connected, Loopback13
L     13.0.0.1/32 is directly connected, Loopback13
O E2  14.0.0.0/8 [110/20] via 192.168.103.30, 00:08:46, GigabitEthernet0/0
O IA  192.168.1.0/24 [110/3] via 192.168.103.30, 00:21:39, GigabitEthernet0/0
O IA  192.168.2.0/24 [110/4] via 192.168.103.30, 00:21:39, GigabitEthernet0/0
O IA  192.168.11.0/24 [110/3] via 192.168.103.30, 00:21:39, GigabitEthernet0/0
O IA  192.168.12.0/24 [110/6] via 192.168.103.30, 00:21:39, GigabitEthernet0/0
O IA  192.168.21.0/24 [110/5] via 192.168.103.30, 00:21:39, GigabitEthernet0/0
O IA  192.168.22.0/24 [110/4] via 192.168.103.30, 00:21:39, GigabitEthernet0/0
O IA  192.168.31.0/24 [110/2] via 192.168.103.30, 00:21:39, GigabitEthernet0/0
O IA  192.168.32.0/24 [110/4] via 192.168.103.30, 00:21:39, GigabitEthernet0/0
O IA  192.168.41.0/24 [110/5] via 192.168.103.30, 00:21:39, GigabitEthernet0/0
O IA  192.168.42.0/24 [110/4] via 192.168.103.30, 00:21:39, GigabitEthernet0/0
192.168.103.0/24 is variably subnetted, 2 subnets, 2 masks
C     192.168.103.0/24 is directly connected, GigabitEthernet0/0
L     192.168.103.31/32 is directly connected, GigabitEthernet0/0
O IA  192.168.104.0/24
       [110/5] via 192.168.103.30, 00:21:39, GigabitEthernet0/0
O IA  192.168.134.0/24
       [110/11] via 192.168.103.30, 00:21:39, GigabitEthernet0/0
```

- Step 2** Configure an OSPF area between the CE and PE routers in your pod to be the stub.
- Step 3** On the CE router, examine the IP routing table, and you should see the intra-area and interarea OSPF routes. There are no external OSPF routes in the routing table; instead, you should see a default route. Outputs should be similar to the following:

```
CE3# show ip route
```

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

```
Gateway of last resort is 192.168.103.30 to network 0.0.0.0
```

```
O*IA 0.0.0.0/0 [110/2] via 192.168.103.30, 00:05:56, GigabitEthernet0/0
      10.0.0.0/8 is variably subnetted, 8 subnets, 2 masks
O IA   10.0.1.1/32 [110/3] via 192.168.103.30, 00:05:56, GigabitEthernet0/0
O IA   10.0.2.1/32 [110/4] via 192.168.103.30, 00:05:56, GigabitEthernet0/0
O IA   10.1.1.1/32 [110/4] via 192.168.103.30, 00:05:56, GigabitEthernet0/0
O IA   10.2.1.1/32 [110/5] via 192.168.103.30, 00:05:56, GigabitEthernet0/0
O IA   10.3.1.1/32 [110/2] via 192.168.103.30, 00:05:56, GigabitEthernet0/0
C      10.3.10.1/32 is directly connected, Loopback0
O IA   10.4.1.1/32 [110/5] via 192.168.103.30, 00:05:56, GigabitEthernet0/0
O IA   10.4.10.0/24 [110/6] via 192.168.103.30, 00:05:56, GigabitEthernet0/0
      13.0.0.0/8 is variably subnetted, 3 subnets, 3 masks
O      13.0.0.0/8 is a summary, 00:17:41, Null0
C      13.0.0.0/24 is directly connected, Loopback13
L      13.0.0.1/32 is directly connected, Loopback13
O IA   192.168.1.0/24 [110/3] via 192.168.103.30, 00:05:56, GigabitEthernet0/0
O IA   192.168.2.0/24 [110/4] via 192.168.103.30, 00:05:56, GigabitEthernet0/0
O IA   192.168.11.0/24 [110/3] via 192.168.103.30, 00:05:56, GigabitEthernet0/0
O IA   192.168.12.0/24 [110/6] via 192.168.103.30, 00:05:56, GigabitEthernet0/0
O IA   192.168.21.0/24 [110/5] via 192.168.103.30, 00:05:56, GigabitEthernet0/0
O IA   192.168.22.0/24 [110/4] via 192.168.103.30, 00:05:56, GigabitEthernet0/0
O IA   192.168.31.0/24 [110/2] via 192.168.103.30, 00:05:56, GigabitEthernet0/0
O IA   192.168.32.0/24 [110/4] via 192.168.103.30, 00:05:56, GigabitEthernet0/0
O IA   192.168.41.0/24 [110/5] via 192.168.103.30, 00:05:56, GigabitEthernet0/0
O IA   192.168.42.0/24 [110/4] via 192.168.103.30, 00:05:56, GigabitEthernet0/0
      192.168.103.0/24 is variably subnetted, 2 subnets, 2 masks
C      192.168.103.0/24 is directly connected, GigabitEthernet0/0
L      192.168.103.31/32 is directly connected, GigabitEthernet0/0
O IA   192.168.104.0/24
      [110/5] via 192.168.103.30, 00:05:56, GigabitEthernet0/0
O IA   192.168.134.0/24
      [110/11] via 192.168.103.30, 00:05:56, GigabitEthernet0/0
```

- Step 4** Configure an OSPF area between the CE and PE routers in your pod to be totally stubby.
- Step 5** On the CE router, examine the IP routing table, and you should see only the intra-area OSPF routes. There are no interarea or external OSPF routes in the routing table; instead, you should see a default route. Output should be similar to the following:

```
CE3# show ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       + - replicated route, % - next hop override
```

Gateway of last resort is 192.168.103.30 to network 0.0.0.0

```
O*IA 0.0.0.0/0 [110/2] via 192.168.103.30, 00:02:03, GigabitEthernet0/0
      10.0.0.0/32 is subnetted, 1 subnets
C      10.3.10.1 is directly connected, Loopback0
      13.0.0.0/8 is variably subnetted, 3 subnets, 3 masks
O      13.0.0.0/8 is a summary, 00:21:54, Null0
C      13.0.0.0/24 is directly connected, Loopback13
L      13.0.0.1/32 is directly connected, Loopback13
      192.168.103.0/24 is variably subnetted, 2 subnets, 2 masks
C      192.168.103.0/24 is directly connected, GigabitEthernet0/0
L      192.168.103.31/32 is directly connected, GigabitEthernet0/0
```

- Step 6** On the PE router, adjust the OSPF cost of the injected default route to 10.
- Step 7** Examine the OSPF database on the CE router. There should be a summary LSA that includes a default route with metric 10. Output should be similar to the following:

```
CE3# show ip ospf database summary 0.0.0.0

      OSPF Router with ID (30.30.30.30) (Process ID 1)

      Summary Net Link States (Area 13)

Routing Bit Set on this LSA in topology Base with MTID 0
LS age: 72
Options: (No TOS-capability, DC, Upward)
LS Type: Summary Links (Network)
Link State ID: 0.0.0.0 (summary Network Number)
Advertising Router: 3.3.3.3
LS Seq Number: 80000003
Checksum: 0xAD79
Length: 28
Network Mask: /0
      MTID: 0      Metric: 10
```

- Step 8** Examine the OSPF default route metric on the CE router. The metric of the default route is 11. Outputs should be similar to the following:

```
CE3# show ip route 0.0.0.0
Routing entry for 0.0.0.0/0, supernet
  Known via "ospf 1", distance 110, metric 11, candidate default path, type inter-area
  Last update from 192.168.103.30 on GigabitEthernet0/0, 00:01:50 ago
  Routing Descriptor Blocks:
    * 192.168.103.30, from 3.3.3.3, 00:01:50 ago, via GigabitEthernet0/0
      Route metric is 11, traffic share count is 1
```

- Step 9** Because the OSPF area between the CE and PE routers is configured as the stub, there is no IP reachability to Loopback1x or Loopback1y (where x or y is your pod number). The stub area does not support an ASBR router inside the area. To get IP reachability, change the stub area into a totally NSSA. A totally NSSA allows an ASBR router inside the area.

Verification

You have completed this task when you attain these results:

- Step 1** On the PE router in your pod, verify the OSPF database. There should be a type 7 LSA that is originated by the CE router. A type 7 LSA describes the CE Loopback1x or Loopback1y network (where x or y is your pod number). Output should be similar to the following:

```
RP/0/RSP0/CPU0:PE3# show ospf database nssa-external
Fri May 16 12:49:18.763 UTC

    OSPF Router with ID (3.3.3.3) (Process ID 1)

      Type-7 AS External Link States (Area 13)

Routing Bit Set on this LSA
LS age: 108
Options: (No TOS-capability, Type 7/5 translation, DC)
LS Type: AS External Link
Link State ID: 13.0.0.0 (External Network Number)
Advertising Router: 30.30.30.30
LS Seq Number: 80000001
Checksum: 0x45b3
Length: 36
Network Mask: /8
    Metric Type: 2 (Larger than any link state path)
    TOS: 0
    Metric: 20
    Forward Address: 10.3.10.1
    External Route Tag: 0
```

- Step 2** The PE router translates the type 7 LSA into a type 5 LSA. On the PE router in your pod, verify the OSPF database. There should be a type 5 LSA that is originated by the PE router. The type 5 LSA describes the CE Loopback1x or Loopback1y network (where x or y is your pod number). Outputs should be similar to the following:

```
RP/0/RSP0/CPU0:PE3# show ospf database external 13.0.0.0
Fri May 16 12:52:11.140 UTC
```

```
OSPF Router with ID (3.3.3.3) (Process ID 1)
```

```
Type-5 AS External Link States
```

```
LS age: 1934
Options: (No TOS-capability, DC)
LS Type: AS External Link
Link State ID: 13.0.0.0 (External Network Number)
Advertising Router: 30.30.30.30
LS Seq Number: 80000003
Checksum: 0xb266
Length: 36
Network Mask: /8
    Metric Type: 2 (Larger than any link state path)
    TOS: 0
    Metric: 20
    Forward Address: 0.0.0.0
    External Route Tag: 0
```

- Step 3** Verify IP connectivity from the CE router in your pod to the Loopback1y or Loopback1x IP address on the neighbor CE router (same team) (where y or x is neighbor pod number). The ping should be successful.

```
CE3# ping 14.0.0.1 (Lo14 IP address from CE4 router)
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 14.0.0.1, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

Hardware Lab 3: Implement Integrated IS-IS Routing

Activity Objective

In this lab activity, you will configure Integrated IS-IS routing by enabling IS-IS to run in Level 1 and Level 2.


Note	Students from two different pods are working as 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 running Cisco IOS XR Software, and the second pod in the same team will work on the PE router running Cisco IOS XE Software. Students in the same team should coordinate their lab activity.
-------------	---

In the lab activity, you will work on different Cisco routers 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 routers to be in a specific IS-IS area
- Enable IS-IS on the correct interfaces
- Configure a router to be a Level 1 IS-IS router
- Configure routers to perform route summarization
- Enable IPv6 IS-IS routing

Visual Objective

The figure illustrates what you will accomplish in this activity.



Required Resources

These are the resources and equipment that 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/IOS XE Commands

Command	Description
configure terminal	Enters configuration mode.
interface <i>interface</i>	Enters interface configuration mode.
ip address <i>ip_address subnet_mask</i>	Sets a primary or secondary IPv4 address for an interface and the subnet mask.
ip ipv6 router isis	Enables IS-IS routing to the interface.
ipv6 address <i>ip_address/subnet_mask</i>	Sets the IPv6 address for an interface and the subnet mask.
ipv6 enable	Enables IPv6 support on the interface.
ipv6 unicast-routing	Enables IPv6 unicast routing (global).
isis circuit-type <i>level-1 level-1-2 level-2-only</i>	Enables the IS-IS routing process to establish the selected IS-IS circuit-type on the interface.
is-type <i>level-1 level-1-2 level-2-only</i>	Sets the IS-IS router to act only in the selected level. The default is acting in both levels.
net <i>net_address</i>	Sets the NET address to the IS-IS router.
ping <i>dest_IP source source_IP</i>	Verifies connectivity between the source IP and the destination IP.
router isis	Creates an IS-IS process.
show ip interface brief	Displays the interface status and the IPv4 addresses that are configured.
show ip route	Displays the current routes in the routing table.
show isis database	Displays the content of the IS-IS database.
show isis database <i>LSP_ID</i> detail	Displays the additional information about a specific LSP ID.
show isis neighbors	Displays IS-IS neighbor information.
show isis topology	Displays lists of information that are related to the IS-IS topology for a specific router.

Command	Description
show running-config	Displays the running configuration.
summary-address <i>network subnet_mask level-1 level-1-2 level-2</i>	Summarizes the IS-IS route to the selected level. The default is summarizing into Level 2.

Cisco IOS XR Commands

Command	Description
address-family <i>ipv4 ipv6 unicast</i>	Enables IPv4 or IPv6 IS-IS routing and enters address family configuration mode for IS-IS (in IS-IS configuration mode)
circuit-type <i>level-1 level-1-2 level-2-only</i>	Enables the IS-IS routing process to establish the selected IS-IS circuit-type on the interface
commit	Commits changes to the running configuration
configure terminal	Enters configuration mode
interface <i>interface (global)</i>	Enters interface configuration mode
interface <i>interface (router)</i>	Defines the interfaces on which the IS-IS protocol runs
ipv6 address <i>ip_address/len</i>	Sets the IPv6 address for an interface and the subnet mask using the prefix length format
ipv6 enable	Enables IPv6 support on the interface
net <i>net_address</i>	Sets the NET address to the IS-IS router
ping <i>dest_IP source source_IP</i>	Verifies connectivity between the source IP and the destination IP (IPv4 and IPv6)
router isis <i>process-ID</i>	Creates an IS-IS process
show ipv4 interface brief	Displays the interface status and IPv4 addresses that are configured
show isis database	Displays the content of the IS-IS database
show isis database <i>LSP_ID detail</i>	Displays the additional information about a specific LSP ID.
show isis neighbors	Displays IS-IS neighbor information
show isis topology	Displays lists of information related to the IS-IS topology for a specific router
show route	Displays the current routes in the routing table

Command	Description
show running-config	Displays the running configuration
single-topology	Enables IS-IS to support a single topology. Used in the IPv6 IS-IS address family.
summary-prefix network/subnet_mask level-1 level-1-2 level-2	Summarizes the IS-IS route to the selected level. The default is summarizing into Level 2.

Task 1: Enabling Integrated IS-IS on the Routers

In this task, you will enable Integrated IS-IS. Pod routers will be configured in the correct IS-IS area, and interfaces will be enabled for IS-IS routing. Additionally, you will configure the CE router to be a Level 1 IS-IS router.

Outputs in this task are taken from team 2 with pods 3 and 4.

Activity Procedure

Complete these steps:

- Step 1** On the CE and PE routers in your team, verify that the Gigabit Ethernet and Loopback0 interfaces are up and running and that correct IPv4 addresses are assigned to the interfaces. Find the correct IP addresses in the visual objective figure.

On the CE router, there should first be Gigabit Ethernet and Loopback interfaces up and running with assigned IP addresses.

```
CE3# show ip interface brief
```

Interface	IP-Address	OK?	Method	Status	Protocol
Embedded-Service-Engine0/0	unassigned	YES	NVRAM	administratively down	down
GigabitEthernet0/0	192.168.103.31	YES	NVRAM	up	up
GigabitEthernet0/1	unassigned	YES	NVRAM	down	down
Loopback0	10.3.10.1	YES	NVRAM	up	up
Loopback13	13.0.0.1	YES	NVRAM	up	up

On the PE router, there should be first and second Gigabit Ethernet and Loopback interfaces up and running with assigned IP addresses.

```
RP/0/RSP0/CPU0:PE3# show ipv4 interface brief
```

```
Mon May 26 10:58:56.303 UTC
```

Interface	IP-Address	Status	Protocol
Loopback0	10.3.1.1	Up	Up
MgmtEth0/RSP0/CPU0/0	10.10.10.25	Up	Up
MgmtEth0/RSP0/CPU0/1	unassigned	Shutdown	Down
GigabitEthernet0/0/0/0	192.168.103.30	Up	Up
GigabitEthernet0/0/0/1	192.168.134.30	Up	Up
GigabitEthernet0/0/0/2	192.168.31.30	Up	Up
GigabitEthernet0/0/0/3	192.168.32.30	Up	Up

```
<... output omitted ...>
```

Step 2 You should be able to successfully test IP connectivity between the CE and PE routers in your pod and between the PE routers in your team.

The ping is successful from the PE to CE router.

```
RP/0/RSP0/CPU0:PE3# ping 192.168.103.31
Mon May 26 11:01:10.483 UTC
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.103.31, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

The ping is successful between the PE routers in your team.

```
RP/0/RSP0/CPU0:PE3# ping 192.168.134.40
Mon May 26 11:02:17.882 UTC
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.134.40, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

In the following steps, you will configure basic IS-IS routing in your team routers:

Step 3 Define IS-IS areas, prepare a NET addressing plan for routers, and determine interfaces where IS-IS will be configured. Use the visual objective and NET addressing sections to complete this step:

Router	IS-IS Area	System ID+00	IS-IS Interfaces
CE (your pod)			
PE (your pod)			
CE (neighbor pod)			
PE (neighbor pod)			

Step 4 On the CE and PE routers in your pod, disable OSPF and enable IS-IS routing, configure the NET address, and enable Integrated IS-IS on the interfaces.

Step 5 Configure the CE router to be a Level 1 IS-IS router only, and configure the PE router to establish only Level 1 IS-IS adjacency with the CE router and only Level 2 IS-IS adjacency with the PE router in the same team.

Verification

You have completed this task when you attain these results:

- Step 1** On the CE router in your pod, verify the IS-IS neighbors, IS-IS topology, and IS-IS database. There should be one neighbor with Level 1 IS-IS adjacency. The IS-IS topology should show the CE and PE routers. In the IS-IS database, there should be three LSPs. Outputs should be similar to the following:

```
CE3# show isis neighbors
```

System Id	Type	Interface	IP Address	State	Holdtime	Circuit Id
PE3	L1	Gi0/0	192.168.103.30	UP	7	PE3.03

```
CE3# show isis topology
```

IS-IS TID 0 paths to level-1 routers

System Id	Metric	Next-Hop	Interface	SNPA
PE3	10	PE3	Gi0/0	6c9c.ed24.1c28
CE3	--			

```
CE3# show isis database
```

IS-IS Level-1 Link State Database:

LSPID	LSP Seq Num	LSP Checksum	LSP Holdtime	ATT/P/OL
PE3.00-00	0x00000008	0x0204	919	1/0/0
PE3.03-00	0x00000003	0x9F0B	742	0/0/0
CE3.00-00	* 0x00000005	0xD04C	745	0/0/0

```
CE4# show isis neighbors
```

System Id	Type	Interface	IP Address	State	Holdtime	Circuit Id
PE4	L1	Gi0/0	192.168.104.40	UP	26	CE4.01

```
CE4# show isis topology
```

IS-IS TID 0 paths to level-1 routers

System Id	Metric	Next-Hop	Interface	SNPA
PE4	10	PE4	Gi0/0	442b.03e7.9380
CE4	--			

```
CE4# show isis database
```

IS-IS Level-1 Link State Database:

LSPID	LSP Seq Num	LSP Checksum	LSP Holdtime	ATT/P/OL
PE4.00-00	0x00000062	0xA104	735	1/0/0
CE4.00-00	* 0x0000005D	0xBFEO	392	0/0/0
CE4.01-00	* 0x00000058	0xB329	702	0/0/0

Step 2 On the PE router in your pod, verify the IS-IS neighbors, IS-IS topology, IS-IS database, and IS-IS database for the CE router in your pod. There should be one neighbor with Level 1 IS-IS adjacency and one neighbor with Level 2 IS-IS adjacency. The IS-IS topology should show the CE and PE routers in the Level 1 topology and the CE and two PE routers in the Level 2 topology. The IS-IS database should have entries in the Level 1 and Level 2 topologies. The IS-IS database for the CE router should have an entry for the remote network of the CE Loopback interface with metric 10 because this is the directly connected network. Outputs should be similar to the following:

RP/0/RSP0/CPU0:PE3# show isis neighbors

Mon May 26 12:34:42.110 UTC

IS-IS 1 neighbors:

System Id	Interface	SNPA	State	Holdtime	Type	IETF-NSF
CE3	Gi0/0/0/0	2c54.2d9b.d3e0	Up	27	L1	Capable
PE4	Gi0/0/0/1	442b.03e7.9381	Up	24	L2	Capable

Total neighbor count: 2

RP/0/RSP0/CPU0:PE3# show isis topology

Mon May 26 12:35:22.793 UTC

IS-IS 1 paths to IPv4 Unicast (Level-1) routers

System Id	Metric	Next-Hop	Interface	SNPA
PE3	--			
CE3	10	CE3	Gi0/0/0/0	2c54.2d9b.d3e0

IS-IS 1 paths to IPv4 Unicast (Level-2) routers

System Id	Metric	Next-Hop	Interface	SNPA
PE3	--			
PE4	10	PE4	Gi0/0/0/1	442b.03e7.9381

RP/0/RSP0/CPU0:PE3# show isis database

Mon May 26 12:36:27.567 UTC

IS-IS 1 (Level-1) Link State Database

LSPID		LSP Seq Num	LSP Checksum	LSP Holdtime	ATT/P/OL
PE3.00-00	*	0x0000000b	0xfb07	1046	1/0/0
PE3.03-00		0x00000006	0x990e	773	0/0/0
CE3.00-00		0x00000008	0xca4f	966	0/0/0

Total Level-1 LSP count: 3 Local Level-1 LSP count: 1

IS-IS 1 (Level-2) Link State Database

LSPID		LSP Seq Num	LSP Checksum	LSP Holdtime	ATT/P/OL
PE3.00-00	*	0x0000000c	0xfbb	801	0/0/0
PE3.01-00		0x00000005	0x2a70	1033	0/0/0
PE4.00-00		0x0000000b	0xe162	1068	0/0/0

Total Level-2 LSP count: 3 Local Level-2 LSP count: 1

RP/0/RSP0/CPU0:PE3# show isis database CE3.00-00 detail

Mon May 26 12:37:31.432 UTC

IS-IS 1 (Level-1) Link State Database

LSPID		LSP Seq Num	LSP Checksum	LSP Holdtime	ATT/P/OL
CE3.00-00		0x0000005d	0x3761	799	0/0/0

Area Address: 49.0003

NLPID: 0xcc

NLPID: 0x8e

Hostname: CE3

IP Address: 10.3.10.1

Metric: 10 IP 10.3.10.1/32

Metric: 10 IP 192.168.103.0/24

IS-IS 1 (Level-2) Link State Database

LSPID		LSP Seq Num	LSP Checksum	LSP Holdtime	ATT/P/OL
-------	--	-------------	--------------	--------------	----------

The requested LSP CE3.00-00 was not found in the IS-IS 1 Level-2 LSP Database

```
PE4# show isis neighbors
```

System Id	Type	Interface	IP Address	State	Holdtime	Circuit Id
CE4	L1	Gi0/0/0	192.168.104.41	UP	7	CE4.01
PE3	L2	Gi0/0/1	192.168.134.30	UP	8	PE3.01

```
PE4# show isis topology
```

```
IS-IS TID 0 paths to level-1 routers
```

System Id	Metric	Next-Hop	Interface	SNPA
PE4	--			
CE4	10	CE4	Gi0/0/0	442b.03f2.ba00

```
IS-IS TID 0 paths to level-2 routers
```

System Id	Metric	Next-Hop	Interface	SNPA
PE3	10	PE3	Gi0/0/1	6c9c.ed24.1c29
PE4				

```
PE4# show isis database
```

```
IS-IS Level-1 Link State Database:
```

LSPID	LSP Seq Num	LSP Checksum	LSP Holdtime	ATT/P/OL
PE4.00-00	* 0x00000063	0x9F05	581	1/0/0
CE4.00-00	0x0000005F	0xBBE2	1169	0/0/0
CE4.01-00	0x00000059	0xB12A	621	0/0/0

```
IS-IS Level-2 Link State Database:
```

LSPID	LSP Seq Num	LSP Checksum	LSP Holdtime	ATT/P/OL
PE3.00-00	0x00000066	0x4FDB	883	0/0/0
PE3.01-00	0x0000005A	0x7FC5	431	0/0/0
PE4.00-00	* 0x00000063	0x1016	609	0/0/0

```
PE4# show isis database CE4.00-00 detail
```

```
IS-IS Level-1 LSP CE4.00-00
```

LSPID	LSP Seq Num	LSP Checksum	LSP Holdtime	ATT/P/OL
CE4.00-00	0x0000005F	0xBBE2	1139	0/0/0

```
Area Address: 49.0004
NLPID: 0xCC 0x8E
Hostname: CE4
IP Address: 10.4.10.1
Metric: 10 IP 10.4.10.1 255.255.255.255
Metric: 10 IP 192.168.104.0 255.255.255.0
```

Step 3 On the CE and PE routers, verify that IS-IS routes are present in the IP routing table. Output on the CE router should be similar to the following:

```
CE3# show ip route isis
<... output omitted ...>
Gateway of last resort is 192.168.103.30 to network 0.0.0.0

i*L1 0.0.0.0/0 [115/10] via 192.168.103.30, 00:44:00, GigabitEthernet0/0
    10.0.0.0/32 is subnetted, 2 subnets
i L1 10.3.1.1 [115/20] via 192.168.103.30, 00:44:00, GigabitEthernet0/0
```

```
CE4# show ip route isis
<... output omitted ...>
Gateway of last resort is 192.168.104.40 to network 0.0.0.0

i*L1 0.0.0.0/0 [115/10] via 192.168.104.40, 17:43:34, GigabitEthernet0/0
    10.0.0.0/32 is subnetted, 2 subnets
i L1 10.4.1.1 [115/20] via 192.168.104.40, 17:43:34, GigabitEthernet0/0
```

Output on the PE router should be similar to the following:

```
RP/0/RSP0/CPU0:PE3# show ip route isis
Mon May 26 12:38:12.643 UTC

i L1 10.3.10.1/32 [115/20] via 192.168.103.31, 00:44:46, GigabitEthernet0/0/0/0
i L2 10.4.1.1/32 [115/20] via 192.168.134.40, 00:55:59, GigabitEthernet0/0/0/1
i L2 10.4.10.1/32 [115/30] via 192.168.134.40, 00:44:24, GigabitEthernet0/0/0/1
i L2 192.168.104.0/24 [115/20] via 192.168.134.40, 00:55:59, GigabitEthernet0/0/0/1
```

```
PE4# show ip route isis
<... output omitted ...>
Gateway of last resort is not set

    10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
i L2 10.3.1.1/32
    [115/20] via 192.168.134.30, 17:44:14, GigabitEthernet0/0/1
i L2 10.3.10.0/24
    [115/30] via 192.168.134.30, 17:44:14, GigabitEthernet0/0/1
i su 10.4.10.0/24 [115/20] via 0.0.0.0, 17:44:14, Null0
i L1 10.4.10.1/32
    [115/20] via 192.168.104.41, 17:44:14, GigabitEthernet0/0/0
i L2 192.168.103.0/24
    [115/20] via 192.168.134.30, 17:44:14, GigabitEthernet0/0/1
```

Step 4 Verify IPv4 connectivity between the CE routers in your team. A ping from the CE router in your pod to the CE router Loopback interface in the other pod (same team) should be successful.

```
CE3# ping 10.4.10.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.4.10.1, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

Task 2: IS-IS Route Summarization

In this task, you will configure IS-IS route summarization.

Activity Procedure

Complete these steps:

- Step 1** On the PE router in your pod, summarize the IS-IS route from Level 1 into Level 2. You will summarize route 10.x.10.1/32 into 10.x.10.0/24 or 10.y.10.1/32 into 10.y.10.0/24, where x or y is your pod number.

Verification

You have completed this task when you attain these results:

- Step 1** On the PE router in your pod, verify that the router generates an IS-IS summary route in the routing table.

```
RP/0/RSP0/CPU0:PE3# show route isis
Mon May 26 12:48:46.207 UTC
```

```
i su 10.3.10.0/24 [115/20] via 0.0.0.0, 00:01:36, Null0
i L1 10.3.10.1/32 [115/20] via 192.168.103.31, 00:55:20, GigabitEthernet0/0/0/0
i L2 10.4.1.1/32 [115/20] via 192.168.134.40, 01:06:32, GigabitEthernet0/0/0/1
i L2 10.4.10.0/24 [115/30] via 192.168.134.40, 00:01:13, GigabitEthernet0/0/0/1
i L2 192.168.104.0/24 [115/20] via 192.168.134.40, 01:06:32, GigabitEthernet0/0/0/1
```

Output taken from the PE router running Cisco IOS XE should be similar to the following:

```
PE4# show ip route isis
<... output omitted ...>
Gateway of last resort is not set

10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
i L2 10.3.1.1/32
    [115/20] via 192.168.134.30, 01:07:06, GigabitEthernet0/0/1
i L2 10.3.10.0/24
    [115/30] via 192.168.134.30, 00:02:07, GigabitEthernet0/0/1
i su 10.4.10.0/24 [115/20] via 0.0.0.0, 00:01:46, Null0
i L1 10.4.10.1/32
    [115/20] via 192.168.104.41, 00:55:31, GigabitEthernet0/0/0
i L2 192.168.103.0/24
    [115/20] via 192.168.134.30, 01:07:06, GigabitEthernet0/0/1
```

- Step 2** On the PE router in your pod, verify that the router is receiving a summary route from other PE routers in the same team.

Output taken from the PE router running Cisco IOS XR should be similar to the following:

```
RP/0/RSP0/CPU0:PE3# show route isis
Mon May 26 12:48:46.207 UTC
```

```
i su 10.3.10.0/24 [115/20] via 0.0.0.0, 00:01:36, Null0
i L1 10.3.10.1/32 [115/20] via 192.168.103.31, 00:55:20, GigabitEthernet0/0/0/0
i L2 10.4.1.1/32 [115/20] via 192.168.134.40, 01:06:32, GigabitEthernet0/0/0/1
i L2 10.4.10.0/24 [115/30] via 192.168.134.40, 00:01:13, GigabitEthernet0/0/0/1
i L2 192.168.104.0/24 [115/20] via 192.168.134.40, 01:06:32, GigabitEthernet0/0/0/1
```

Output taken from the PE router running Cisco IOS XE should be similar to the following:

```

PE4# show ip route isis
<... output omitted ...>
Gateway of last resort is not set

    10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
i L2    10.3.1.1/32
        [115/20] via 192.168.134.30, 01:08:36, GigabitEthernet0/0/1
i L2    10.3.10.0/24
        [115/30] via 192.168.134.30, 00:03:37, GigabitEthernet0/0/1
i su    10.4.10.0/24 [115/20] via 0.0.0.0, 00:03:16, Null0
i L1    10.4.10.1/32
        [115/20] via 192.168.104.41, 00:57:01, GigabitEthernet0/0/0
i L2    192.168.103.0/24
        [115/20] via 192.168.134.30, 01:08:36, GigabitEthernet0/0/1

```

- Step 3** Verify IPv4 connectivity between the CE routers in your team. A ping from the CE router in your pod to the CE router Loopback interface in the other pod (same team) should still be successful.

```

CE3# ping 10.4.10.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.4.10.1, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms

```

Task 3: Enabling IPv6 IS-IS Routing

In this task, you will enable Integrated IS-IS for IPv6.

Activity Procedure

Complete these steps to verify that routers in your team are correctly preconfigured:

- Step 1** On the CE and PE routers in your team, verify that IPv6 routing is enabled on the Gigabit Ethernet and Loopback0 interfaces and that the IPv6 address is assigned to the CE and PE router Loopback interfaces. Find the correct IP addresses in the visual objective or IP addressing sections.

If IPv6 is not preconfigured on the CE router, enter commands similar to these (x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned):

```

CEx(config)#
ipv6 unicast-routing
!
interface Loopback0
ipv6 address 2001:DB8:10:x:10::1/128
ipv6 enable
!
interface GigabitEthernet0/0
ipv6 enable

```

```

CEy(config)#
ipv6 unicast-routing
!
interface Loopback0
  ipv6 address 2001:DB8:10:y:10::1/128
  ipv6 enable
!
interface GigabitEthernet0/0
  ipv6 enable

```

If IPv6 is not preconfigured on the PE router, enter commands similar to these (x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned):

```

RP/0/RSP0/CPU0:PEx(config)#
interface Loopback0
  ipv6 address 2001:db8:10:1:1::1/128
  ipv6 enable
!
interface GigabitEthernet0/0/0/0
  ipv6 enable
!
interface GigabitEthernet0/0/0/1
  ipv6 enable
!
commit

```

```

PEy(config)#
ipv6 unicast-routing
!
interface Loopback0
  ipv6 address 2001:DB8:10:2:1::1/128
  ipv6 enable
!
interface GigabitEthernet0/0/0
  ipv6 enable
!
interface GigabitEthernet0/0/1
  ipv6 enable

```

Step 2 On the CE and PE routers in your pod, disable OSPFv3 and enable IPv6 IS-IS routing. On the CE router, enable IPv6 IS-IS on the first Gigabit Ethernet and Loopback interfaces. On the PE router, enable IPv6 IS-IS on the first and second Gigabit Ethernet and Loopback interfaces.

Verification

You have completed this task when you attain these results:

Step 1 On the CE and PE routers, verify that IPv6 IS-IS routes are present in the IPv6 routing table. Output on the CE router should be similar to the following:

```
CE3# show ipv6 route isis
IPv6 Routing Table - default - 4 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
        B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
        H - NHRP, I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
        IS - ISIS summary, D - EIGRP, EX - EIGRP external, NM - NEMO
        ND - ND Default, NDp - ND Prefix, DCE - Destination, NDr - Redirect
        O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
        ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, ls - LISP site
        ld - LISP dyn-EID, a - Application

I1 ::/0 [115/10]
    via FE80::6E9C:EDFF:FE24:1C28, GigabitEthernet0/0
I1 2001:DB8:10:3:1::1/128 [115/20]
    via FE80::6E9C:EDFF:FE24:1C28, GigabitEthernet0/0
```

```
CE4# show ipv6 route isis
IPv6 Routing Table - default - 4 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
        B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
        H - NHRP, I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
        IS - ISIS summary, D - EIGRP, EX - EIGRP external, NM - NEMO
        ND - ND Default, NDp - ND Prefix, DCE - Destination, NDr - Redirect
        O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
        ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, ls - LISP site
        ld - LISP dyn-EID, a - Application

I1 ::/0 [115/10]
    via FE80::462B:3FF:FEE7:9380, GigabitEthernet0/0
I1 2001:DB8:10:4:1::1/128 [115/20]
    via FE80::462B:3FF:FEE7:9380, GigabitEthernet0/0
```

Output on the PE router should be similar to the following:

```
RP/0/RSP0/CPU0:PE3# show route ipv6 isis
Mon May 26 13:30:45.532 UTC

i L1 2001:db8:10:3:10::1/128
    [115/20] via fe80::2e54:2dff:fe9b:d3e0, 00:07:29, GigabitEthernet0/0/0/0
i L2 2001:db8:10:4:1::1/128
    [115/20] via fe80::462b:3ff:fee7:9381, 00:05:41, GigabitEthernet0/0/0/1
i L2 2001:db8:10:4:10::1/128
    [115/30] via fe80::462b:3ff:fee7:9381, 00:05:37, GigabitEthernet0/0/0/1
```

```

PE4# show ipv6 route isis
IPv6 Routing Table - default - 5 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
        B - BGP, R - RIP, H - NHRP, I1 - ISIS L1
        I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary, D - EIGRP
        EX - EIGRP external, ND - ND Default, NDp - ND Prefix, DCE - Destination
        NDr - Redirect, O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1
        OE2 - OSPF ext 2, ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
        ls - LISP site, ld - LISP dyn-EID, a - Application
I2 2001:DB8:10:3:1::1/128 [115/20]
    via FE80::6E9C:EDFF:FE24:1C29, GigabitEthernet0/0/1
I2 2001:DB8:10:3:10::1/128 [115/30]
    via FE80::6E9C:EDFF:FE24:1C29, GigabitEthernet0/0/1
I1 2001:DB8:10:4:10::1/128 [115/20]
    via FE80::462B:3FF:FEF2:BA00, GigabitEthernet0/0/0

```

Step 2 Verify IPv6 connectivity between the CE routers in your team. A ping from the CE router in your pod to the CE router Loopback interface in the other pod (same team) should be successful.

```

CE3# ping 2001:DB8:10:4:10::1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:10:4:10::1, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms

```

Hardware Lab 4: Implement Basic BGP Routing

Activity Objective

In this lab activity, you will configure internal and external BGP routing.


Students from two different pods are working as 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 running Cisco IOS XR Software, and the second pod in the same team will work on the PE router running Cisco IOS XE Software. Students in the same team should coordinate their lab activity.

In the lab activity, you will work on different Cisco routers 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 EBGP between pod routers
- Configure IBGP between the pod router and the backbone router
- Configure BGP authentication
- Configure BGP next-hop self

Visual Objective

The figure illustrates what you will accomplish in this activity.



Required Resources

These are the resources and equipment that 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 the interface on the router
cdp enable	Enables Cisco Discovery Protocol on the interface
configure terminal	Enters configuration mode
interface <i>interface</i>	Enters interface configuration mode
ip address <i>ip_address subnet_mask</i>	Sets a primary or secondary IPv4 address for an interface and the subnet mask
ip ipv6 router isis	Enables IS-IS routing to the interface
ipv6 enable	Enables IPv6 support on the interface
ipv6 unicast-routing	Enables IPv6 unicast routing (global)
isis circuit-type <i>level-1 level-1-2 level-2-only</i>	Enables the IS-IS routing process to establish the selected IS-IS circuit-type on the interface
neighbor <i>IP-address next-hop-self</i>	Enables the internal BGP router to send BGP routes with the router BGP IP address
neighbor <i>IP-address password password</i>	Enables authentication for the BGP neighbor
neighbor <i>IP-address remote-as AS-number</i>	Configures the BGP neighbor
neighbor <i>IP-address update-source interface</i>	Enables the BGP router to send BGP packets from the source IP address
network <i>network mask subnet-mask</i>	Enables the BGP router to include the network from the routing table into BGP routing
ping <i>dest_IP source source_IP</i>	Verifies connectivity between the source IP and the destination IP
router bgp <i>AS-number</i>	Creates a BGP process and enters BGP process configuration mode
router isis	Creates an IS-IS process and enters IS-IS process configuration mode
show ip bgp	Displays the BGP routing table

Command	Description
show ip bgp summary	Displays BGP routing protocol characteristics, including the BGP neighbor status
show ip interface brief	Displays the interface status and IPv4 addresses that are configured
show ip route	Displays the current routes in the routing table
show isis database	Displays the content of the IS-IS database
show isis neighbors	Displays IS-IS neighbor information
show isis topology	Displays lists of information that are related to the IS-IS topology for a specific router
show running-config	Displays the running configuration

Cisco IOS XR Commands

Command	Description
[no] shutdown	Enables or disables the interface on the router
address-family ipv4 ipv6 unicast	Enables IPv4 or IPv6 IS-IS or BGP routing and enters address family configuration mode for IS-IS or BGP (in router IS-IS or BGP configuration mode)
circuit-type level-1 level-1-2 level-2-only	Enables the IS-IS routing process to establish the selected IS-IS circuit-type on the interface
commit	Commits changes to the running configuration
configure terminal	Enters configuration mode
interface <i>interface</i> (global)	Enters interface configuration mode
interface <i>interface</i> (router)	Defines the interfaces on which the IS-IS protocol runs
ip address <i>ip_address mask</i>	Sets the IPv4 address for an interface
ipv6 enable	Enables IPv6 support on the interface
neighbor <i>IP-address</i>	Configures the BGP neighbor and enters BGP neighbor configuration mode
network <i>network/subnet-mask</i>	Enables the BGP router to include the network from the routing table into BGP routing (BGP address family mode)

Command	Description
next-hop-self	Enables the internal BGP router to send BGP routes with the router BGP IP address (BGP neighbor address family mode)
pass	Passes the route for further processing (route-policy configuration mode)
password clear <i>password</i>	Enables authentication for the BGP neighbor
ping <i>dest_IP source source_IP</i>	Verifies connectivity between the source IP and the destination IP (IPv4 and IPv6)
remote-as <i>AS-number</i>	Configures the AS number for the BGP neighbor (BGP neighbor mode)
route-policy <i>route-policy-name</i>	Creates the route policy and enters route policy configuration mode
route-policy <i>route-policy-name</i> in out	Applies the route policy to the BGP neighbor
router bgp <i>AS-number</i>	Creates a BGP process and enters BGP process configuration mode
router isis <i>process-ID</i>	Creates an IS-IS process
show bgp	Displays the BGP routing table
show bgp summary	Displays BGP routing protocol characteristics, including the BGP neighbor status
show ipv4 interface brief	Displays the interface status and IPv4 addresses that are configured
show isis database	Displays the content of the IS-IS database
show isis neighbors	Displays IS-IS neighbor information
show isis topology	Displays lists of information that are related to the IS-IS topology for a specific router
show route	Displays the current routes in the routing table
show running-config	Displays the running configuration
update-source <i>interface</i>	Enables the BGP router to send BGP packets from the source IP address (BGP neighbor configuration mode)

Task 1: Configuring External BGP

In this task, you will enable a secure external BGP between the CE and PE routers. Additionally, you will advertise the network into the BGP process.

Outputs in this task are taken from team 2 with pods 3 and 4.

Activity Procedure

Complete these steps:

- Step 1** On the PE router in your pod, enable the BGP process in the AS 64500 and configure an external BGP neighbor. The external BGP neighbor is the CE router in your pod with the IP address that is configured on the Gi0/0 interface and autonomous system number 6450x or 6450y, where x or y is your pod number.
- Step 2** On the CE router in your pod, enable the BGP process in the AS 6450x or AS 6450y, where x or y is your pod number, and configure the external BGP neighbor. The external BGP neighbor is the PE router in your pod with the IP address that is configured on the first Gigabit Ethernet interface and autonomous system number 64500.
- Step 3** Verify that the BGP session is established between the PE and CE routers in your pod. In the “state/prefix received” column, you should see the number zero.

```
RP/0/RSP0/CPU0:PE3# show bgp summary
```

```
<... output omitted ...>
```

Neighbor	Spk	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	St/PfxRcd
192.168.103.31	0	64503	80	74	2	0	0	00:53:36	0!

```
PE4# show ip bgp summary
```

```
BGP router identifier 10.4.1.1, 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
192.168.104.41	4	64504	70	71	1	0	0	01:00:51	0

```
CE3# show ip bgp summary
```

```
BGP router identifier 13.0.0.1, local AS number 64503
```

```
BGP table version is 1, main routing table version 1
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
192.168.103.30	4	64500	60	68	1	0	0	00:57:56	0

```
CE4# show ip bgp summary
```

```
BGP router identifier 14.0.0.1, local AS number 64504
```

```
BGP table version is 1, main routing table version 1
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
192.168.104.40	4	64500	73	72	1	0	0	01:03:23	0

- Step 4** Configure BGP authentication between the PE and CE routers in your pod. Use the password "cisco."
- Step 5** On the CE and PE routers in your pod, advertise the Loopback0 network into BGP.
- Step 6** On the PE (Cisco IOS XR only) router, make sure that the router will accept and advertise all BGP routes.

Verification

You have completed this task when you attain these results:

- Step 1** Verify that the BGP session is established between the PE and CE routers in your pod. In the "state/prefix received" column, you should see number "1."

```
RP/0/RSP0/CPU0:PE3# show bgp summary
Tue May 27 12:24:58.110 UTC
BGP router identifier 10.3.1.1, local AS number 64500
BGP generic scan interval 60 secs
BGP table state: Active
Table ID: 0xe0000000 RD version: 4
BGP main routing table version 4
BGP scan interval 60 secs
```

BGP is operating in STANDALONE mode.

Process	RcvTblVer	bRIB/RIB	LabelVer	ImportVer	SendTblVer	StandbyVer
Speaker	4	4	4	4	4	4

Neighbor	Spk	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	St/PfxRcd
192.168.103.31	0	64503	164	154	4	0	0	00:53:51	1

```
PE4# show ip bgp summary
BGP router identifier 10.4.1.1, local AS number 64500
BGP table version is 3, main routing table version 3
2 network entries using 496 bytes of memory
2 path entries using 240 bytes of memory
2/2 BGP path/bestpath attribute entries using 480 bytes of memory
1 BGP AS-PATH entries using 24 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 1240 total bytes of memory
BGP activity 2/0 prefixes, 2/0 paths, scan interval 60 secs
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
192.168.104.41	4	64504	175	175	3	0	0	02:35:36	1

CE3# show ip bgp summary

BGP router identifier 13.0.0.1, local AS number 64503
BGP table version is 3, main routing table version 3
2 network entries using 288 bytes of memory
2 path entries using 160 bytes of memory
2/2 BGP path/bestpath attribute entries using 288 bytes of memory
1 BGP AS-PATH entries using 24 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 760 total bytes of memory
BGP activity 2/0 prefixes, 2/0 paths, scan interval 60 secs

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
192.168.103.30	4	64500	81	89	3	0	0	01:17:08	i

CE4# show ip bgp summary

BGP router identifier 14.0.0.1, local AS number 64504
BGP table version is 3, main routing table version 3
2 network entries using 288 bytes of memory
2 path entries using 160 bytes of memory
2/2 BGP path/bestpath attribute entries using 288 bytes of memory
1 BGP AS-PATH entries using 24 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 760 total bytes of memory
BGP activity 2/0 prefixes, 2/0 paths, scan interval 60 secs

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
192.168.104.40	4	64500	176	176	3	0	0	02:37:06	i

Step 2 Verify that the BGP routing table appears on the PE and CE routers in your pod. You should see the BGP route from the neighbor.

CE3# show ip bgp

BGP table version is 3, local router ID is 13.0.0.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

	Network	Next Hop	Metric	LocPrf	Weight	Path
*>	10.3.1.1/32	192.168.103.30	0		0	64500 i
*>	10.3.10.1/32	0.0.0.0	0		32768	i

```
RP/0/RSP0/CPU0:PE3# show bgp
Tue May 27 12:56:26.065 UTC
BGP router identifier 10.3.1.1, local AS number 64500
BGP generic scan interval 60 secs
BGP table state: Active
Table ID: 0xe0000000 RD version: 4
BGP main routing table version 4
BGP scan interval 60 secs

Status codes: s suppressed, d damped, h history, * valid, > best
               i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network        Next Hop           Metric LocPrf Weight Path
*> 10.3.1.1/32     0.0.0.0                 0           32768 i
*> 10.3.10.1/32    192.168.103.31          0           0 64503 i

Processed 2 prefixes, 2 paths
```

```
CE4# show ip bgp
BGP table version is 3, local router ID is 14.0.0.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
               x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

   Network        Next Hop           Metric LocPrf Weight Path
*> 10.4.1.1/32     192.168.104.40          0           0 64500 i
*> 10.4.10.1/32    0.0.0.0                0           32768 i
```

```
PE4# show ip bgp
BGP table version is 3, local router ID is 10.4.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
               x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

   Network        Next Hop           Metric LocPrf Weight Path
*> 10.4.1.1/32     0.0.0.0                 0           32768 i
*> 10.4.10.1/32    192.168.104.41          0           0 64504 i
```

Task 2: Configuring Internal BGP

In this task, you will enable an internal BGP between the PE router in your pod and the P1 router in the backbone. Additionally, you will enable the BGP next-hop self feature on the PE router.

Activity Procedure

Complete these steps:

- Step 1** On the PE router, enable the interface and IS-IS Level-2 routing toward the P1 router for the IPv4 and IPv6 address families.
- Step 2** On the PE router in your pod, configure an internal BGP neighbor. The internal BGP neighbor is the P1 router in the backbone with IP address 10.0.1.1. The PE router in your pod should source BGP packets from the Loopback0 IP address.
- The P1 router in the backbone is configured as the BGP route reflector.
- Step 3** Configure the PE router in your pod to set the Loopback0 interface IP address for all BGP routes that are sent to the P1 BGP neighbor.

Verification

You have completed this task when you attain these results:

- Step 1** On the PE router in your pod, you should see BGP routes from the neighbor pod, with next-hop IP addresses that are set to the neighbor PE router Loopback0 interface IP address. The neighbor pod PE router should have routes coming from your pod PE router with your PE router Loopback0 interface IP address.

```
RP/0/RSP0/CPU0:PE3# show bgp
Tue May 27 13:42:12.519 UTC
BGP router identifier 10.3.1.1, local AS number 64500
BGP generic scan interval 60 secs
BGP table state: Active
Table ID: 0xe0000000 RD version: 9
BGP main routing table version 9
BGP scan interval 60 secs

Status codes: s suppressed, d damped, h history, * valid, > best
               i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network        Next Hop        Metric LocPrf Weight Path
*> 10.3.1.1/32     0.0.0.0                0           32768 i
*> 10.3.10.1/32    192.168.103.31         0           0 64503 i
*>i10.4.1.1/32     10.4.1.1               0          100           0 i
*>i10.4.10.1/32    10.4.1.1               0          100           0 64504 i
* i10.5.100.1/32   172.16.5.1             0          1000           0 ?
*>i209.165.201.0/27 10.0.2.1               0          100           0 64497 64498 ?
*>i209.165.202.128/27 10.0.2.1              0          100           0 64498 64495 64496 ?

Processed 7 prefixes, 7 paths
```

```
PE4# show ip bgp
```

```
BGP table version is 8, local router ID is 10.4.1.1
```

```
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,  
r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,  
x best-external, a additional-path, c RIB-compressed,
```

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

```
RPKI validation codes: V valid, I invalid, N Not found
```

	Network	Next Hop	Metric	LocPrf	Weight	Path
r>i	10.3.1.1/32	10.3.1.1	0	100	0	i
*>i	10.3.10.1/32	10.3.1.1	0	100	0	64503 i
*>	10.4.1.1/32	0.0.0.0	0		32768	i
*>	10.4.10.1/32	192.168.104.41	0		0	64504 i
* i	10.5.100.1/32	172.16.5.1	0	1000	0	?
*>i	209.165.201.0/27	10.0.2.1	0	100	0	64497 64498 ?
*>i	209.165.202.128/27	10.0.2.1	0	100	0	64498 64495 64496 ?

Step 2 Verify IP connectivity between the CE routers in your team. A ping from the CE router in your pod to the CE router in the other pod (same team) should be successful.

```
CE3# ping 10.4.10.1
```

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

```
Sending 5, 100-byte ICMP Echos to 10.4.10.1, timeout is 2 seconds:
```

```
!!!!!
```

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


Hardware Lab 5: Implement Route Redistribution

Activity Objective

In this lab activity, you will enable connectivity between two routing domains by configuring route redistribution between IS-IS and OSPF. You will configure two-way redistribution, and you will protect the network from routing information loops. By modifying the administrative distance, you will enable the pod PE router to use the optimal path to the backbone. Finally, you will reduce the number of routing updates that is sent to the pod CE router by configuring one-way redistribution.


Students from two different pods are working as 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 running Cisco IOS XR Software, and the second pod in the same team will work on the PE router running Cisco IOS XE Software. Students in the same team should coordinate their lab activity.

In the lab activity, you will work on different Cisco routers 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 two-way redistribution between IS-IS and OSPF
- Prevent potential routing loops between IS-IS and OSPF routing protocols
- Modify the administrative distance of the OSPF routing protocol
- Configure one-way redistribution between IS-IS and OSPF

Visual Objective

The figure illustrates what you will accomplish in this activity.



Required Resources

These are the resources and equipment that 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] ip ipv6 router isis	Enables or disables IS-IS routing to the interface
[no] isis circuit-type level-1 level-1-2 level-2-only	Enables or disables the IS-IS routing process to establish the selected IS-IS circuit-type on the interface
[no] router isis	Creates or deletes an IS-IS process
configure terminal	Enters configuration mode
default-information originate always	Originates the OSPF default route, even if there is no default route in the routing table (always)
distance ospf external <i>value</i>	Modifies the OSPF external administrative distance
interface <i>interface</i>	Enters interface configuration mode
network <i>prefix wildcard_mask area area</i>	(Router) Enables OSPF routing on the network and places the network into the OSPF area
ping <i>dest_IP source source_IP</i>	Verifies connectivity between the source IP and the destination IP
redistribute <i>protocol</i>	Redistributes protocol into OSPF or IS-IS
router ospf <i>process-id</i>	Enables OSPFv2 processes on the router
set match tag <i>value</i>	Routes the map set or match option to set or match the tag value
show ip interface brief	Displays the interface status and IPv4 addresses that are configured
show ip ospf	Displays information that is related to the OSPF routing protocol running on the router
show ip ospf interface	Displays OSPF interface information
show ip ipv6 ospf database	Displays the content of the OSPF database
show ip ipv6 ospf neighbors	Displays OSPF neighbor information
show ip ipv6 protocols	Displays IPv4 or IPv6 protocols running on the router

Command	Description
show ip ipv6 route	Displays the current routes in the routing table
show isis database	Displays the content of the IS-IS database
show isis neighbors	Displays IS-IS neighbor information
show isis topology	Displays lists of information that is related to the IS-IS topology for a specific router
show running-config	Displays the running configuration

Cisco IOS XR Commands

Command	Description
address-family ipv4 ipv6 unicast	Enables IPv4 or IPv6 IS-IS or OSPF routing and enters address family configuration mode
area area	(Router) Creates the OSPF area on the router
commit	Commits changes to the running configuration
configure terminal	Enters configuration mode
default-information originate always	Originates the OSPF default route, even if there is no default route in the routing table (always)
distance ospf external <i>value</i>	Modifies the OSPF external administrative distance
if <i>condition</i> then <i>action-1</i> else <i>action-2</i> endif	Routes the policy list if-then sentence
interface interface	(Router) Defines the interfaces on which the OSPF or IS-IS protocol runs
pass	Passes the route for further processing (route-policy configuration mode)
ping <i>dest_IP</i> source <i>source_IP</i>	Verifies connectivity between the source IP and the destination IP (IPv4 and IPv6)
redistribute <i>protocol</i>	Redistributes the protocol into OSPF or IS-IS
route-policy <i>route-policy-name</i>	Creates the route policy and enters route policy configuration mode
router isis <i>process-ID</i>	Creates an IS-IS process and enters router IS-IS process configuration mode

Command	Description
router ospf ospfv3 process-ID	Creates an OSPFv2 or OSPFv3 process and enters router OSPFv2 or OSPFv3 process configuration mode
show ipv4 interface brief	Displays the interface status and IPv4 addresses that are configured
show isis database	Displays the content of the IS-IS database
show isis neighbors	Displays IS-IS neighbor information
show isis topology	Displays lists of information that are related to the IS-IS topology for a specific router
show ospf	Displays information that is related to the OSPF routing protocol running on the router
show ospf ospfv3 database	Displays the content of the OSPF database
show ospf ospfv3 interface	Displays OSPF interface information
show ospf ospfv3 neighbors	Displays OSPF neighbor information
show protocols	Displays the protocols running on the router
show route	Displays the current routes in the routing table
show running-config	Displays the running configuration

Task 1: Redistributing between OSPF and IS-IS

In this task, you will configure two-way redistribution by redistributing OSPF routes into the IS-IS routing protocol and IS-IS routes into the OSPF routing protocol. Redistribution will be performed on the PE router.


Outputs in this task are taken from team 2 with pods 3 and 4.

Activity Procedure

Prepare the initial routing protocol setup for redistribution:

- Step 1** On your pod CE router, disable the IS-IS routing protocol and enable OSPF process 1. Start OSPF Area 1x or Area 1y on the first Gigabit Ethernet and Loopback0 interfaces, where x or y are your pod numbers.
- Step 2** On your pod PE router, disable IS-IS routing on the first and second Gigabit Ethernet interfaces. Enable OSPF process 1. Start OSPF Area 1x or Area 1y (where x or y are your pod numbers) on the first Gigabit Ethernet interface and Area 0 on the second Gigabit Ethernet interface.

The figure illustrates what you have accomplished in previous steps. There are IS-IS and OSPF routing protocols running on your pod devices. This routing setup will be used in the following route redistribution steps.



- Step 3** On the pod routers, verify that the IS-IS and OSPF neighbors are up and verify that the pod CE router receives OSPF routes from other pods in the same team. From the pod CE router, check that you are not able to ping the core devices.

On the PE router, the **show** command output should be similar to the following:

```
RP/0/RSP0/CPU0:PE3# show isis neighbors
```

```
Mon Jun  2 13:24:35.427 UTC
```

```
IS-IS 1 neighbors:
```

System Id	Interface	SNPA	State	Holdtime	Type	IETF-NSF
P1	Gi0/0/0/2	6c9c.ed24.19ab	Up	29	L2	Capable

```
Total neighbor count: 1
```

```
RP/0/RSP0/CPU0:PE3# show ospf neighbor
```

```
Mon Jun  2 13:42:52.626 UTC
```

```
* Indicates MADJ interface
```

```
Neighbors for OSPF 1
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
10.4.1.1	1	FULL/DR	00:00:39	192.168.134.40	GigabitEthernet0/0/0/1
Neighbor is up for 00:05:27					
10.3.10.1	0	FULL/DROTHER	00:00:31	192.168.103.31	GigabitEthernet0/0/0/0
Neighbor is up for 00:00:55					

```
Total neighbor count: 2
```

```
PE4# show isis neighbors
```

```
Tag null:
```

System Id	Type	Interface	IP Address	State	Holdtime	Circuit Id
P1	L2	Gi0/0/2	192.168.41.1	UP	9	P1.01

```
PE4# show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
10.3.1.1	1	FULL/BDR	00:00:33	192.168.134.30	GigabitEthernet0/0/1
14.0.0.1	0	FULL/DROTHER	00:00:35	192.168.104.41	GigabitEthernet0/0/0

On the CE router, the **show** command output should be similar to the following. There should be a summary LSA for Loopback0 from the neighbor pod CE router that is present in the OSPF database:

```
CE3# show ip ospf database
```

```
OSPF Router with ID (10.3.10.1) (Process ID 1)
```

```
Router Link States (Area 13)
```

Link ID	ADV Router	Age	Seq#	Checksum	Link count
10.3.1.1	10.3.1.1	877	0x80000004	0x001E15	1
10.3.10.1	10.3.10.1	876	0x80000002	0x0041B8	2

```
Net Link States (Area 13)
```

Link ID	ADV Router	Age	Seq#	Checksum
192.168.103.30	10.3.1.1	877	0x80000001	0x0064B0

```
Summary Net Link States (Area 13)
```

Link ID	ADV Router	Age	Seq#	Checksum
10.4.10.1	10.3.1.1	65	0x80000001	0x00FC15
192.168.104.0	10.3.1.1	1146	0x80000001	0x00EF6A
192.168.134.0	10.3.1.1	1148	0x80000001	0x009AA2

```
CE4# show ip ospf database
```

```
OSPF Router with ID (10.4.10.1) (Process ID 1)
```

```
Router Link States (Area 14)
```

Link ID	ADV Router	Age	Seq#	Checksum	Link count
10.4.1.1	10.4.1.1	135	0x80000006	0x000D0C	1
10.4.10.1	10.4.10.1	133	0x80000002	0x00459B	2

```
Net Link States (Area 14)
```

Link ID	ADV Router	Age	Seq#	Checksum
192.168.104.40	10.4.1.1	135	0x80000001	0x000106

```
Summary Net Link States (Area 14)
```

Link ID	ADV Router	Age	Seq#	Checksum
10.3.10.1	10.4.1.1	941	0x80000001	0x005BAD
192.168.103.0	10.4.1.1	1203	0x80000005	0x004508
192.168.134.0	10.4.1.1	547	0x80000002	0x00EA47

A ping from the CE router to the backbone, P1 Loopback0, should *not* be successful:

```
CE3# ping 10.0.1.1
```

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

```
Sending 5, 100-byte ICMP Echos to 10.0.1.1, timeout is 2 seconds:
```

```
.....
```

```
Success rate is 0 percent (0/5)
```

- Step 4** On the pod PE router, mutually redistribute between the IS-IS and OSPF routing protocols. Redistribute only IS-IS Level 2 routes into OSPF and make OSPF routes redistribute into IS-IS as Level 2.

Verification

You have completed this task when you attain these results:

- Step 1** On the pod PE router on the OSPF database, you should see routes that are redistributed from the IS-IS routing protocol as external routes.

```
RP/0/RSP0/CPU0:PE3# show ospf database
```

```
<... output omitted ...>
```

Type-5 AS External Link States

Link ID	ADV Router	Age	Seq#	Checksum	Tag
10.0.1.1	10.3.1.1	235	0x80000001	0x00d0b4	0
10.0.2.1	10.3.1.1	235	0x80000001	0x00c5be	0
10.3.1.1	10.3.1.1	235	0x80000001	0x00acd5	0
10.4.1.1	10.3.1.1	235	0x80000001	0x00a0e0	0
192.168.1.0	10.3.1.1	235	0x80000001	0x00ab7b	0
192.168.2.0	10.3.1.1	235	0x80000001	0x00a085	0
192.168.11.0	10.3.1.1	235	0x80000001	0x003ddf	0
192.168.12.0	10.3.1.1	235	0x80000001	0x0032e9	0
192.168.21.0	10.3.1.1	235	0x80000001	0x00ce44	0
192.168.22.0	10.3.1.1	235	0x80000001	0x00c34e	0
192.168.31.0	10.3.1.1	235	0x80000001	0x0060a8	0
192.168.32.0	10.4.1.1	222	0x80000001	0x004db9	0
192.168.41.0	10.3.1.1	235	0x80000001	0x00f10d	0
192.168.42.0	10.3.1.1	235	0x80000001	0x00e617	0

```
PE4# show ip ospf database
```

```
<... output omitted ...>
```

Type-5 AS External Link States

Link ID	ADV Router	Age	Seq#	Checksum	Tag
10.0.1.1	10.3.1.1	301	0x80000001	0x00D0B4	0
10.0.2.1	10.3.1.1	301	0x80000001	0x00C5BE	0
10.3.1.1	10.3.1.1	301	0x80000001	0x00ACD5	0
10.4.1.1	10.3.1.1	301	0x80000001	0x00A0E0	0
192.168.1.0	10.3.1.1	301	0x80000001	0x00AB7B	0
192.168.2.0	10.3.1.1	301	0x80000001	0x00A085	0
192.168.11.0	10.3.1.1	301	0x80000001	0x003DDF	0
192.168.12.0	10.3.1.1	301	0x80000001	0x0032E9	0
192.168.21.0	10.3.1.1	301	0x80000001	0x00CE44	0
192.168.22.0	10.3.1.1	301	0x80000001	0x00C34E	0
192.168.31.0	10.3.1.1	301	0x80000001	0x0060A8	0
192.168.32.0	10.4.1.1	286	0x80000001	0x004DB9	0
192.168.41.0	10.3.1.1	301	0x80000001	0x00F10D	0
192.168.42.0	10.3.1.1	301	0x80000001	0x00E617	0

Task 2: Preventing Potential Routing Loops


In this task, you will prevent potential routing loops between IS-IS and OSPF routing protocols. Because two-way redistribution is configured between IS-IS and OSPF on two routers, there is a possibility for routing information loops. You will make redistribution more predictable.

Activity Procedure

Complete these steps:

- Step 1** On the pod PE router, configure a Route Policy Language (RPL) configuration (Cisco IOS XR) or route map (Cisco IOS XE) configuration to deny routes with tag 115. Permit other routes and set tag 110. Use an RPL or route map name of OSPFintoISIS.
- Step 2** On the pod PE router, apply the RPL or route map that is configured in the previous step to the OSPF into the IS-IS redistribution configuration.
- Step 3** On the pod PE router, configure the RPL (Cisco IOS XR) configuration or route map (Cisco IOS XE) configuration to deny routes with tag 110, permit other routes, and set tag 115. Use an RPL or route map name of ISISintoOSPF.
- Step 4** On the pod PE router, apply the RPL or route map that is configured in the previous step to the IS-IS into the OSPF redistribution configuration.

The figure illustrates what you have accomplished in previous steps. There is a route-tagging and filtering mechanism that is used when two-way redistribution is configured on two routers.



Verification

You have completed this task when you attain these results:

Step 1 On the pod PE router on the OSPF database, you should see routes that are redistributed from the IS-IS routing protocol as external routes with tag 115.

```
RP/0/RSP0/CPU0:PE3# show ospf database
```

```
<... output omitted ...>
```

```
Type-5 AS External Link States
```

Link ID	ADV Router	Age	Seq#	Checksum	Tag
10.0.1.1	10.3.1.1	994	0x8000001e	0x00b440	115
10.0.2.1	10.3.1.1	994	0x8000001e	0x00a94a	115
10.3.1.1	10.3.1.1	994	0x8000001e	0x009061	115
10.4.1.1	10.3.1.1	994	0x8000001e	0x00846c	115
192.168.1.0	10.3.1.1	994	0x8000001e	0x008f07	115
192.168.2.0	10.3.1.1	994	0x8000001e	0x008411	115
192.168.11.0	10.3.1.1	994	0x8000001e	0x00216b	115
192.168.12.0	10.3.1.1	994	0x8000001e	0x001675	115
192.168.21.0	10.3.1.1	994	0x8000001e	0x00b2cf	115
192.168.22.0	10.3.1.1	994	0x8000001e	0x00a7d9	115
192.168.31.0	10.3.1.1	994	0x8000001e	0x004434	115
192.168.32.0	10.4.1.1	897	0x80000001	0x006b28	115
192.168.41.0	10.3.1.1	994	0x8000001e	0x00d598	115
192.168.42.0	10.3.1.1	994	0x8000001e	0x00caa2	115

```
PE4# show ip ospf database
```

```
<... output omitted ...>
```

```
Type-5 AS External Link States
```

Link ID	ADV Router	Age	Seq#	Checksum	Tag
10.0.1.1	10.3.1.1	1126	0x8000001E	0x00B440	115
10.0.2.1	10.3.1.1	1126	0x8000001E	0x00A94A	115
10.3.1.1	10.3.1.1	1126	0x8000001E	0x009061	115
10.4.1.1	10.3.1.1	1126	0x8000001E	0x00846C	115
192.168.1.0	10.3.1.1	1126	0x8000001E	0x008F07	115
192.168.2.0	10.3.1.1	1126	0x8000001E	0x008411	115
192.168.11.0	10.3.1.1	1126	0x8000001E	0x00216B	115
192.168.12.0	10.3.1.1	1126	0x8000001E	0x001675	115
192.168.21.0	10.3.1.1	1126	0x8000001E	0x00B2CF	115
192.168.22.0	10.3.1.1	1126	0x8000001E	0x00A7D9	115
192.168.31.0	10.3.1.1	1126	0x8000001E	0x004434	115
192.168.32.0	10.4.1.1	1026	0x80000001	0x006B28	115
192.168.41.0	10.3.1.1	1126	0x8000001E	0x00D598	115
192.168.42.0	10.3.1.1	1126	0x8000001E	0x00CAA2	115

Step 2 On the pod PE router, verify the routing table. On one PE router in the team, there are IS-IS routes, but the other PE router in the team is showing the same routes as OSPF external routes. This is not optimal routing because one PE router is using the PE-to-PE link to reach networks in the backbone.

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

```
<... output omitted ...>
```

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

```
i L2 10.0.1.1/32 [115/20] via 192.168.31.1, 00:21:11, GigabitEthernet0/0/0/2
i L2 10.0.2.1/32 [115/30] via 192.168.31.1, 00:21:11, GigabitEthernet0/0/0/2
L 10.3.1.1/32 is directly connected, 2w5d, Loopback0
B 10.3.10.1/32 [20/0] via 192.168.103.31, 16:30:53
i L2 10.4.1.1/32 [115/30] via 192.168.31.1, 17:02:07, GigabitEthernet0/0/0/2
O IA 10.4.10.1/32 [110/3] via 192.168.134.40, 16:17:28, GigabitEthernet0/0/0/1
C 10.10.10.0/24 is directly connected, 3w0d, MgmtEth0/RSP0/CPU0/0
L 10.10.10.25/32 is directly connected, 3w0d, MgmtEth0/RSP0/CPU0/0
i L2 192.168.1.0/24 [115/20] via 192.168.31.1, 00:21:11, GigabitEthernet0/0/0/2
i L2 192.168.2.0/24 [115/20] via 192.168.31.1, 00:21:11, GigabitEthernet0/0/0/2
i L2 192.168.11.0/24 [115/20] via 192.168.31.1, 00:21:11, GigabitEthernet0/0/0/2
i L2 192.168.12.0/24 [115/30] via 192.168.31.1, 00:21:11, GigabitEthernet0/0/0/2
i L2 192.168.21.0/24 [115/20] via 192.168.31.1, 00:21:11, GigabitEthernet0/0/0/2
i L2 192.168.22.0/24 [115/30] via 192.168.31.1, 00:21:11, GigabitEthernet0/0/0/2
C 192.168.31.0/24 is directly connected, 2w5d, GigabitEthernet0/0/0/2
L 192.168.31.30/32 is directly connected, 2w5d, GigabitEthernet0/0/0/2
C 192.168.32.0/24 is directly connected, 2w5d, GigabitEthernet0/0/0/3
L 192.168.32.30/32 is directly connected, 2w5d, GigabitEthernet0/0/0/3
i L2 192.168.41.0/24 [115/20] via 192.168.31.1, 17:02:07, GigabitEthernet0/0/0/2
i L2 192.168.42.0/24 [115/30] via 192.168.31.1, 6d17h, GigabitEthernet0/0/0/2
C 192.168.103.0/24 is directly connected, 2w4d, GigabitEthernet0/0/0/0
L 192.168.103.30/32 is directly connected, 2w4d, GigabitEthernet0/0/0/0
O IA 192.168.104.0/24 [110/2] via 192.168.134.40, 16:35:30, GigabitEthernet0/0/0/1
C 192.168.134.0/24 is directly connected, 2w0d, GigabitEthernet0/0/0/1
L 192.168.134.30/32 is directly connected, 2w0d, GigabitEthernet0/0/0/1
B 209.165.201.0/27 [200/0] via 10.0.2.1, 6d16h
B 209.165.202.128/27 [200/0] via 10.0.2.1, 6d16h
```

```

PE4# show ip route
<... output omitted ...>
Gateway of last resort is not set

10.0.0.0/32 is subnetted, 6 subnets
O E2 10.0.1.1 [110/20] via 192.168.134.30, 00:20:29, GigabitEthernet0/0/1
O E2 10.0.2.1 [110/20] via 192.168.134.30, 00:20:29, GigabitEthernet0/0/1
O E2 10.3.1.1 [110/20] via 192.168.134.30, 00:20:29, GigabitEthernet0/0/1
O IA 10.3.10.1 [110/12] via 192.168.134.30, 16:30:16, GigabitEthernet0/0/1
C 10.4.1.1 is directly connected, Loopback0
B 10.4.10.1 [20/0] via 192.168.104.41, 16:16:41
O E2 192.168.1.0/24
[110/20] via 192.168.134.30, 00:20:29, GigabitEthernet0/0/1
O E2 192.168.2.0/24
[110/20] via 192.168.134.30, 00:20:29, GigabitEthernet0/0/1
O E2 192.168.11.0/24
[110/20] via 192.168.134.30, 00:20:29, GigabitEthernet0/0/1
O E2 192.168.12.0/24
[110/20] via 192.168.134.30, 00:20:29, GigabitEthernet0/0/1
O E2 192.168.21.0/24
[110/20] via 192.168.134.30, 00:20:29, GigabitEthernet0/0/1
O E2 192.168.22.0/24
[110/20] via 192.168.134.30, 00:20:29, GigabitEthernet0/0/1
O E2 192.168.31.0/24
[110/20] via 192.168.134.30, 00:20:29, GigabitEthernet0/0/1
i L2 192.168.32.0/24 [115/30] via 192.168.41.1, 6d17h, GigabitEthernet0/0/2
192.168.41.0/24 is variably subnetted, 2 subnets, 2 masks
C 192.168.41.0/24 is directly connected, GigabitEthernet0/0/2
L 192.168.41.40/32 is directly connected, GigabitEthernet0/0/2
192.168.42.0/24 is variably subnetted, 2 subnets, 2 masks
C 192.168.42.0/24 is directly connected, GigabitEthernet0/0/3
L 192.168.42.40/32 is directly connected, GigabitEthernet0/0/3
O IA 192.168.103.0/24
[110/11] via 192.168.134.30, 16:34:38, GigabitEthernet0/0/1
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
192.168.134.0/24 is variably subnetted, 2 subnets, 2 masks
C 192.168.134.0/24 is directly connected, GigabitEthernet0/0/1
L 192.168.134.40/32 is directly connected, GigabitEthernet0/0/1
209.165.201.0/27 is subnetted, 1 subnets
B 209.165.201.0 [200/0] via 10.0.2.1, 6d16h
209.165.202.0/27 is subnetted, 1 subnets
B 209.165.202.128 [200/0] via 10.0.2.1, 6d16h

```

Task 3: Modifying Administrative Distance

In this task, you will modify the administrative distance (AD) of the OSPF routing protocol to make the OSPF external routes less preferable than IS-IS routes. This change will make the team PE routers use the optimal path to the backbone networks.

Activity Procedure

Complete this step:

Step 1 On the pod PE router, modify the AD for external OSPF routes to a value of one greater than the IS-IS AD.

Verification

You have completed this task when you attain these results:

Step 1 On the pod PE router, verify the routing table. Both PE routers in the team will show the IS-IS routes for backbone networks, representing optimal routing.

```
RP/0/RSP0/CPU0:PE3# show route
<... output omitted ...>
Gateway of last resort is not set
```

```
i L2 10.0.1.1/32 [115/20] via 192.168.31.1, 01:07:25, GigabitEthernet0/0/0/2
i L2 10.0.2.1/32 [115/30] via 192.168.31.1, 01:07:25, GigabitEthernet0/0/0/2
L 10.3.1.1/32 is directly connected, 2w5d, Loopback0
B 10.3.10.1/32 [20/0] via 192.168.103.31, 17:17:07
i L2 10.4.1.1/32 [115/30] via 192.168.31.1, 17:48:22, GigabitEthernet0/0/0/2
O IA 10.4.10.1/32 [110/3] via 192.168.134.40, 00:11:20, GigabitEthernet0/0/0/1
C 10.10.10.0/24 is directly connected, 3w0d, MgmtEth0/RSP0/CPU0/0
L 10.10.10.25/32 is directly connected, 3w0d, MgmtEth0/RSP0/CPU0/0
i L2 192.168.1.0/24 [115/20] via 192.168.31.1, 01:07:25, GigabitEthernet0/0/0/2
i L2 192.168.2.0/24 [115/20] via 192.168.31.1, 01:07:25, GigabitEthernet0/0/0/2
i L2 192.168.11.0/24 [115/20] via 192.168.31.1, 01:07:25, GigabitEthernet0/0/0/2
i L2 192.168.12.0/24 [115/30] via 192.168.31.1, 01:07:25, GigabitEthernet0/0/0/2
i L2 192.168.21.0/24 [115/20] via 192.168.31.1, 01:07:25, GigabitEthernet0/0/0/2
i L2 192.168.22.0/24 [115/30] via 192.168.31.1, 01:07:25, GigabitEthernet0/0/0/2
C 192.168.31.0/24 is directly connected, 2w5d, GigabitEthernet0/0/0/2
L 192.168.31.30/32 is directly connected, 2w5d, GigabitEthernet0/0/0/2
C 192.168.32.0/24 is directly connected, 2w5d, GigabitEthernet0/0/0/3
L 192.168.32.30/32 is directly connected, 2w5d, GigabitEthernet0/0/0/3
i L2 192.168.41.0/24 [115/20] via 192.168.31.1, 17:48:22, GigabitEthernet0/0/0/2
i L2 192.168.42.0/24 [115/30] via 192.168.31.1, 6d17h, GigabitEthernet0/0/0/2
C 192.168.103.0/24 is directly connected, 2w4d, GigabitEthernet0/0/0/0
L 192.168.103.30/32 is directly connected, 2w4d, GigabitEthernet0/0/0/0
O IA 192.168.104.0/24 [110/2] via 192.168.134.40, 00:11:20, GigabitEthernet0/0/0/1
C 192.168.134.0/24 is directly connected, 2w0d, GigabitEthernet0/0/0/1
L 192.168.134.30/32 is directly connected, 2w0d, GigabitEthernet0/0/0/1
B 209.165.201.0/27 [200/0] via 10.0.2.1, 6d17h
B 209.165.202.128/27 [200/0] via 10.0.2.1, 6d17h
```

```
PE4# show isis route
<... output omitted ...>
```

Gateway of last resort is not set

```
10.0.0.0/32 is subnetted, 6 subnets
i L2 10.0.1.1 [115/20] via 192.168.41.1, 00:12:53, GigabitEthernet0/0/2
i L2 10.0.2.1 [115/30] via 192.168.41.1, 00:12:53, GigabitEthernet0/0/2
i L2 10.3.1.1 [115/30] via 192.168.41.1, 00:12:53, GigabitEthernet0/0/2
O IA 10.3.10.1 [110/12] via 192.168.134.30, 00:12:53, GigabitEthernet0/0/1
C 10.4.1.1 is directly connected, Loopback0
B 10.4.10.1 [20/0] via 192.168.104.41, 17:05:29
i L2 192.168.1.0/24 [115/20] via 192.168.41.1, 00:12:53, GigabitEthernet0/0/2
i L2 192.168.2.0/24 [115/20] via 192.168.41.1, 00:12:53, GigabitEthernet0/0/2
i L2 192.168.11.0/24
[115/20] via 192.168.41.1, 00:12:53, GigabitEthernet0/0/2
i L2 192.168.12.0/24
[115/30] via 192.168.41.1, 00:12:53, GigabitEthernet0/0/2
i L2 192.168.21.0/24
[115/20] via 192.168.41.1, 00:12:53, GigabitEthernet0/0/2
i L2 192.168.22.0/24
[115/30] via 192.168.41.1, 00:12:53, GigabitEthernet0/0/2
i L2 192.168.31.0/24
[115/20] via 192.168.41.1, 00:12:53, GigabitEthernet0/0/2
i L2 192.168.32.0/24 [115/30] via 192.168.41.1, 6d17h, GigabitEthernet0/0/2
192.168.41.0/24 is variably subnetted, 2 subnets, 2 masks
C 192.168.41.0/24 is directly connected, GigabitEthernet0/0/2
L 192.168.41.40/32 is directly connected, GigabitEthernet0/0/2
192.168.42.0/24 is variably subnetted, 2 subnets, 2 masks
C 192.168.42.0/24 is directly connected, GigabitEthernet0/0/3
L 192.168.42.40/32 is directly connected, GigabitEthernet0/0/3
O IA 192.168.103.0/24
[110/11] via 192.168.134.30, 00:12:53, GigabitEthernet0/0/1
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
192.168.134.0/24 is variably subnetted, 2 subnets, 2 masks
C 192.168.134.0/24 is directly connected, GigabitEthernet0/0/1
L 192.168.134.40/32 is directly connected, GigabitEthernet0/0/1
209.165.201.0/27 is subnetted, 1 subnets
B 209.165.201.0 [200/0] via 10.0.2.1, 6d17h
209.165.202.0/27 is subnetted, 1 subnets
B 209.165.202.128 [200/0] via 10.0.2.1, 6d17h
```

Step 2 On the pod CE router, verify the routing table. You should see external OSPF routes for all backbone networks.

```
CE3# show ip route ospf
<... output omitted ...>
Gateway of last resort is not set

10.0.0.0/32 is subnetted, 6 subnets
O E2 10.0.1.1 [110/20] via 192.168.103.30, 01:11:42, GigabitEthernet0/0
O E2 10.0.2.1 [110/20] via 192.168.103.30, 01:11:42, GigabitEthernet0/0
O E2 192.168.1.0/24 [110/20] via 192.168.103.30, 01:11:42, GigabitEthernet0/0
O E2 192.168.2.0/24 [110/20] via 192.168.103.30, 01:11:42, GigabitEthernet0/0
O E2 192.168.11.0/24
    [110/20] via 192.168.103.30, 01:11:42, GigabitEthernet0/0
O E2 192.168.12.0/24
    [110/20] via 192.168.103.30, 01:11:42, GigabitEthernet0/0
O E2 192.168.21.0/24
    [110/20] via 192.168.103.30, 01:11:42, GigabitEthernet0/0
O E2 192.168.22.0/24
    [110/20] via 192.168.103.30, 01:11:42, GigabitEthernet0/0
O E2 192.168.31.0/24
    [110/20] via 192.168.103.30, 01:11:42, GigabitEthernet0/0
O E2 192.168.32.0/24
    [110/20] via 192.168.103.30, 01:09:59, GigabitEthernet0/0
O E2 192.168.41.0/24
    [110/20] via 192.168.103.30, 01:11:42, GigabitEthernet0/0
O E2 192.168.42.0/24
    [110/20] via 192.168.103.30, 01:11:42, GigabitEthernet0/0
O IA 192.168.104.0/24
    [110/3] via 192.168.103.30, 17:21:25, GigabitEthernet0/0
O IA 192.168.134.0/24
    [110/2] via 192.168.103.30, 17:21:25, GigabitEthernet0/0
```

```
CE4# show ip route ospf
<... output omitted ...>
Gateway of last resort is not set

10.0.0.0/32 is subnetted, 6 subnets
O E2 10.0.1.1 [110/20] via 192.168.104.40, 00:17:43, GigabitEthernet0/0
O E2 10.0.2.1 [110/20] via 192.168.104.40, 00:17:43, GigabitEthernet0/0
O E2 192.168.1.0/24 [110/20] via 192.168.104.40, 00:17:43, GigabitEthernet0/0
O E2 192.168.2.0/24 [110/20] via 192.168.104.40, 00:17:43, GigabitEthernet0/0
O E2 192.168.11.0/24
    [110/20] via 192.168.104.40, 00:17:43, GigabitEthernet0/0
O E2 192.168.12.0/24
    [110/20] via 192.168.104.40, 00:17:43, GigabitEthernet0/0
O E2 192.168.21.0/24
    [110/20] via 192.168.104.40, 00:17:43, GigabitEthernet0/0
O E2 192.168.22.0/24
    [110/20] via 192.168.104.40, 00:17:43, GigabitEthernet0/0
O E2 192.168.31.0/24
    [110/20] via 192.168.104.40, 00:17:43, GigabitEthernet0/0
O E2 192.168.32.0/24
    [110/20] via 192.168.104.40, 01:12:24, GigabitEthernet0/0
O E2 192.168.41.0/24
    [110/20] via 192.168.104.40, 01:14:07, GigabitEthernet0/0
O E2 192.168.42.0/24
    [110/20] via 192.168.104.40, 01:14:07, GigabitEthernet0/0
O IA 192.168.103.0/24
    [110/12] via 192.168.104.40, 17:10:28, GigabitEthernet0/0
O IA 192.168.134.0/24
    [110/11] via 192.168.104.40, 17:10:28, GigabitEthernet0/0
```

Task 4: One-Way Redistribution

In this task, you will modify two-way redistribution into one-way redistribution on the PE router. This will reduce routing update information that is sent to the CE routers.

Activity Procedure

Complete these steps:

- Step 1** On the pod PE router, remove IS-IS into OSPF redistribution.
- Step 2** On the pod PE router, originate an OSPF default route. Allow OSPF to originate the default route, even if there is no default route in the routing table.

Verification

You have completed this task when you attain these results:

- Step 1** On the pod CE router on the routing table, you should see the default OSPF route pointing to the pod PE router.

```
CE3# show ip route ospf
<... output omitted ...>
Gateway of last resort is 192.168.103.30 to network 0.0.0.0

O*E2 0.0.0.0/0 [110/1] via 192.168.103.30, 00:03:05, GigabitEthernet0/0
O IA 192.168.104.0/24
      [110/3] via 192.168.103.30, 17:36:43, GigabitEthernet0/0
O IA 192.168.134.0/24
      [110/2] via 192.168.103.30, 17:36:43, GigabitEthernet0/0
```

```
CE4# show ip route ospf
<... output omitted ...>
Gateway of last resort is 192.168.104.40 to network 0.0.0.0

O*E2 0.0.0.0/0 [110/1] via 192.168.104.40, 00:09:23, GigabitEthernet0/0
O IA 192.168.103.0/24
      [110/12] via 192.168.104.40, 17:24:10, GigabitEthernet0/0
O IA 192.168.134.0/24
      [110/11] via 192.168.104.40, 17:24:10, GigabitEthernet0/0
```

- Step 2** On the pod CE router, verify IP connectivity to the backbone. A ping to the P1 router Loopback0 interface (10.0.1.1) should be successful.

```
CE3# ping 10.0.1.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.0.1.1, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```


Hardware Lab 6: Influence BGP Route Selection

Activity Objective

In this lab activity, you will configure BGP to influence route selection by using the weight, local preference, MED, and community attributes in a situation where support for connections to multiple service providers is needed.


Students from two different pods are working as 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 running Cisco IOS XR Software, and the second pod in the same team will work on the PE router running Cisco IOS XE Software. Students in the same team should coordinate their lab activity.

In the lab activity, you will work on different Cisco routers 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 per-neighbor BGP weights
- Configure BGP local preference using RPLs or route-maps
- Configure BGP MED
- Configure service provider BGP community rules and enable route tagging by using BGP communities

Visual Objective

The figure illustrates what you will accomplish in this activity.



Required Resources

These are the resources and equipment that 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] ip ipv6 router isis	Enables or disables IS-IS routing to the interface
[no] isis circuit-type level-1 level-1-2 level-2-only	Enables or disables the IS-IS routing process to establish the selected IS-IS circuit type on the interface
router isis	Creates or enters an IS-IS process
[no] shutdown	Enables or disables the interface on the router
ip bgp-community new-format	Enables the BGP community new format presentation
ip community-list expanded <i>name</i> permit deny <i>value</i>	Creates the community list
ip ipv6 router isis	Enables IS-IS routing to the interface
isis circuit-type level-1 level-1-2 level-2-only	Enables the IS-IS routing process to establish the selected IS-IS circuit type on the interface
neighbor <i>IP-address</i> ebgp-multihop	Configures the EBGp multihop feature for the BGP neighbor
neighbor <i>IP-address</i> remote-as <i>AS-number</i>	Configures the BGP neighbor
neighbor <i>IP-address</i> route-map <i>name</i> in out	Applies the route map to the BGP neighbor
neighbor <i>IP-address</i> send-community both	Enables the router to send standard and extended BGP communities to the BGP neighbor
neighbor <i>IP-address</i> weight <i>value</i>	Configures the BGP weight for routes that are received from the BGP neighbor
route-map <i>name</i>	Creates and enters the route map
router bgp <i>AS-number</i>	Creates a BGP process and enters BGP process configuration mode
router isis	Creates an IS-IS process and enters IS-IS process configuration mode
set match tag local-preference metric community <i>value</i>	Enables the route map set or match option

Command	Description
show ip bgp	Displays the BGP routing table
show ip bgp summary	Displays BGP routing protocol characteristics, including BGP neighbor status
show isis database neighbors topology	Displays the content of the IS-IS database or neighbor information or displays lists of information that are related to the IS-IS topology for a specific router

Cisco IOS XR Commands

Command	Description
[no] shutdown	Enables or disables the interface on the router
address-family ipv4 ipv6 unicast	Enables IPv4 or IPv6 IS-IS or BGP routing and enters address family configuration mode for IS-IS or BGP (in router IS-IS or BGP configuration mode)
circuit-type level-1 level-1-2 level-2-only	Enables the IS-IS routing process to establish the selected IS-IS circuit type on the interface
commit	Commits changes to the running configuration
ebgp-multihop	Configures the EBGp multihop feature for the BGP neighbor
if condition then action-1 else action-2 endif	Routes the policy list if-then sentence
interface interface (router)	Defines the interfaces on which the IS-IS protocol runs
neighbor IP-address	Configures the BGP neighbor and enters BGP neighbor configuration mode
pass	Passes the route for further processing (route-policy configuration mode)
prepend as-path value repets	Enables AS-path prepending in the RPL
remote-as AS-number	Configures the AS number for the BGP neighbor (BGP neighbor mode)
route-policy route-policy-name	Creates the route policy and enters route policy configuration mode
route-policy route-policy-name in out	Applies the route policy to the BGP neighbor
router bgp AS-number	Creates a BGP process and enters BGP process configuration mode


Command	Description
router isis process-ID	Creates an IS-IS process and enters router IS-IS process configuration mode
set local-preference med value	Sets the BGP local preference or MED in the RPL
show bgp	Displays the BGP routing table
show bgp summary	Displays BGP routing protocol characteristics, including BGP neighbor status
show isis database neighbors topology	Displays the content of the IS-IS database or neighbor information or displays lists of information that are related to the IS-IS topology for a specific router
show route	Displays the current routes in the routing table

Task 1: Setting the Second Link between the CE and PE Routers

Task Visual Objective

In this task, you will enable a second link between your pod CE router and the neighbor pod PE routers.

The figure shows what you will accomplish in this task:



Outputs in this task are taken from team 2 with pods 3 and 4.

Activity Procedure

Perform these steps:

- Step 1** Between the pod CE router and the other pod PE router, configure a second link. On the CE router, use the Gigabit Ethernet 0/1 interface, and on the neighbor pod PE router, use the first Gigabit Ethernet subinterface. Use VLAN x0 or y0, where x or y is your pod number. On the new VLAN, use subnet 192.168.20x.0/24 or 192.168.20y.0/24.
- Step 2** On the pod CE and PE routers, disable OSPF routing.

Verification

You have completed this task when you attain this result:

- Step 1** On the pod CE router, test connectivity to the neighbor pod PE router.

```
CE3# ping 192.168.203.40
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.203.40, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

```
CE4# ping 192.168.204.30
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.204.30, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```


Task 2: Configuring BGP Weight and Local Preference

Task Visual Objective

In this task, you will enable a second EBGP session between the pod CE router and the neighbor pod PE router. You will configure the second EBGP session to become the primary session for BGP routes. On the CE router, you will configure BGP weights, and on the PE router, you will configure BGP local preference by using RPLs or route-maps.

The figure shows what you will accomplish in this task:

Configuring BGP Weight and Local Preference



Activity Procedure

Perform these steps:

Step 1 Verify that the BGP session between the pod PE routers and P1 router is established.

```
RP/0/RSP0/CPU0:PE3# show bgp summary
```

<... output omitted ...>

Process	RcvTblVer	bRIB/RIB	LabelVer	ImportVer	SendTblVer	StandbyVer
Speaker	13	13	13	13	13	13

Neighbor	Spk	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	St/PfxRcd
10.0.1.1	0	64500	18443	18442	13	0	0	1w5d	5
192.168.103.31	0	64503	20531	18661	13	0	0	6d19h	1

```
PE4# show ip bgp summary
```

<... output omitted ...>

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
10.0.1.1	4	64500	18445	20295	15	0	0	1w5d	5
192.168.104.41	4	64504	10755	10764	15	0	0	6d18h	1

- Step 2** Between the pod CE router and the other pod PE router, configure a second EBGP session. On this EBGP session, use IP addresses from a second Gigabit Ethernet interface on the CE router and the Gigabit Ethernet subinterface (configured in task 0) on the PE router.
- Step 3** On the pod CE router, use the per-neighbor BGP weight option to configure BGP weight. Configure the CE router to prefer BGP routes that are received from the neighbor pod PE router.
- Step 4** On the pod PE router, use the RPL (Cisco IOS XR) or route-map command (Cisco IOS XE) to configure BGP local preference. The pod PE router should prefer the internal BGP route (route to neighbor pod PE router) toward the pod CE router.
- The BGP routing protocol will not immediately update the routing table. Use the Cisco IOS and IOS XE **clear ip bgp *** **soft** command or Cisco IOS XR **clear bgp *** command to make BGP immediately update the routing table. An alternative solution to the **clear** command is to disable and enable the pod CE router Loopback0 interface.

Verification

You have completed this task when you attain this result:

- Step 1** On the pod CE router, observe BGP routes 192.168.1.0/24 and 192.168.2.0/24, with the next hop to the neighbor pod PE router, as preferred.

```
CE3# show ip bgp
<... output omitted ...>
      Network          Next Hop           Metric LocPrf Weight Path
  *> 10.3.1.1/32        192.168.203.40             0          2 64500 i
  *                   192.168.103.30             0          1 64500 i
  *> 10.3.10.1/32        0.0.0.0                   0         32768 i
  *> 10.4.1.1/32         192.168.203.40             0          2 64500 i
  *                   192.168.103.30             0          1 64500 i
  *> 10.4.10.1/32        192.168.203.40             0          2 64500 64504 i
  *                   192.168.103.30             0          1 64500 64504 i
  *> 192.168.1.0         192.168.203.40             0          2 64500 i
  *                   192.168.103.30             0          1 64500 i
  *> 192.168.2.0         192.168.203.40             0          2 64500 i
  *                   192.168.103.30             0          1 64500 i
<... output omitted ...>
```

```
CE4# show ip bgp
<... output omitted ...>
      Network          Next Hop           Metric LocPrf Weight Path
  *> 10.3.1.1/32        192.168.204.30             0          2 64500 i
  *                   192.168.104.40             0          1 64500 i
  *> 10.3.10.1/32        192.168.204.30             0          2 64500 64503 i
  *                   192.168.104.40             0          1 64500 64503 i
  *> 10.4.1.1/32         192.168.204.30             0          2 64500 i
  *                   192.168.104.40             0          1 64500 i
  *> 10.4.10.1/32        0.0.0.0                   0         32768 i
  *> 192.168.1.0         192.168.204.30             0          2 64500 i
  *                   192.168.104.40             0          1 64500 i
  *> 192.168.2.0         192.168.204.30             0          2 64500 i
  *                   192.168.104.40             0          1 64500 i
<... output omitted ...>
```

Step 2 The pod PE router should show a preference on the internal BGP route toward the pod CE router.

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

```
<... output omitted ...>
```

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 10.3.1.1/32	0.0.0.0	0		32768	i
*>i10.3.10.1/32	10.4.1.1	0	100	0	64503 i
*	192.168.103.31	0	99	0	64503 i
*>i10.4.1.1/32	10.4.1.1	0	100	0	i
*> 10.4.10.1/32	192.168.204.41	0		0	64504 i
* i10.5.100.1/32	172.16.5.1	0	1000	0	?
*>i192.168.1.0/24	10.0.1.1	0	100	0	i
*>i192.168.2.0/24	10.0.1.1	0	100	0	i
*>i209.165.201.0/27	10.0.2.1	0	100	0	64497 64498 ?
*>i209.165.202.128/27	10.0.2.1	0	100	0	64498 64495 64496 ?

```
Processed 9 prefixes, 10 paths
```

```
PE4# show ip bgp
```

```
<... output omitted ...>
```

Network	Next Hop	Metric	LocPrf	Weight	Path
r>i 10.3.1.1/32	10.3.1.1	0	100	0	i
*> 10.3.10.1/32	192.168.203.31	0		0	64503 i
*> 10.4.1.1/32	0.0.0.0	0		32768	i
*	192.168.104.41	0	99	0	64504 i
*>i	10.3.1.1	0	100	0	64504 i
* i 10.5.100.1/32	172.16.5.1	0	1000	0	?
r>i 192.168.1.0	10.0.1.1	0	100	0	i
r>i 192.168.2.0	10.0.1.1	0	100	0	i
*>i 209.165.201.0/27	10.0.2.1	0	100	0	64497 64498 ?
*>i 209.165.202.128/27	10.0.2.1	0	100	0	64498 64495 64496 ?


Task 3: Configuring BGP Multiexit Discriminator

Task Visual Objective

In this task, you will remove BGP local preference configuration from the previous task and configure the pod CE router to send BGP routes that are set with different MED values. You will configure the second EBGP session to be the primary session for incoming and outgoing traffic.

The figure shows what you will accomplish in this task:

Configuring BGP Multiexit Discriminator



Activity Procedure

Perform these steps:

- Step 1** From the pod PE routers, remove the local preference configuration that was created in the previous task.
- Step 2** On the pod PE router, verify that the direct path is selected to the pod CE router. The following outputs are taken from PE routers:

```
RP/0/RSP0/CPU0:PE3# show bgp
<... output omitted ...>
```

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 10.3.1.1/32	0.0.0.0	0		32768	i
*> 10.3.10.1/32	192.168.103.31	0		0	64503 i
*> 10.4.1.1/32	10.4.1.1	0	100	0	i
*> 10.4.10.1/32	192.168.204.41	0		0	64504 i
*> 10.5.100.1/32	172.16.5.1	0	1000	0	?
*> 192.168.1.0/24	10.0.1.1	0	100	0	i
*> 192.168.2.0/24	10.0.1.1	0	100	0	i
*> 209.165.201.0/27	10.0.2.1	0	100	0	64497 64498 ?
*> 209.165.202.128/27	10.0.2.1	0	100	0	64498 64495 64496 ?

Processed 9 prefixes, 9 paths

```

PE4# show ip bgp
<... output omitted ...>

```

Network	Next Hop	Metric	LocPrf	Weight	Path
r>i 10.3.1.1/32	10.3.1.1	0	100	0	i
* i 10.3.10.1/32	10.3.1.1	0	100	0	64503 i
*>	192.168.203.31	0		0	64503 i
*> 10.4.1.1/32	0.0.0.0	0		32768	i
* 10.4.10.1/32	192.168.104.41	0	99	0	64504 i
*>i	10.3.1.1	0	100	0	64504 i
* i 10.5.100.1/32	172.16.5.1	0	1000	0	?
r>i 192.168.1.0	10.0.1.1	0	100	0	i
r>i 192.168.2.0	10.0.1.1	0	100	0	i
*>i 209.165.201.0/27	10.0.2.1	0	100	0	64497 64498 ?
*>i 209.165.202.128/27	10.0.2.1	0	100	0	64498 64495 64496 ?

Step 3 On the pod CE routers, use route maps to set the BGP MED toward the PE routers. Set a higher BGP MED toward the pod PE router and a lower BGP MED toward the neighbor pod PE router.

Verification

You have completed this task when you attain this result:

Step 1 Verify that the pod PE router selects the nondirect path to the pod CE router.

```

RP/0/RSP0/CPU0:PE3# show bgp
<... output omitted ...>

```

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 10.3.1.1/32	0.0.0.0	0		32768	i
*>i 10.3.10.1/32	10.4.1.1	1000	100	0	64503 i
*>	192.168.103.31	1001		0	64503 i
*>i 10.4.1.1/32	10.4.1.1	0	100	0	i
*> 10.4.10.1/32	192.168.204.41	0		0	64504 i
* i 10.5.100.1/32	172.16.5.1	0	1000	0	?
*>i 192.168.1.0/24	10.0.1.1	0	100	0	i
*>i 192.168.2.0/24	10.0.1.1	0	100	0	i
*>i 209.165.201.0/27	10.0.2.1	0	100	0	64497 64498 ?
*>i 209.165.202.128/27	10.0.2.1	0	100	0	64498 64495 64496 ?

Processed 9 prefixes, 10 paths

```

PE4# show ip bgp
<... output omitted ...>

```


Network	Next Hop	Metric	LocPrf	Weight	Path
r>i 10.3.1.1/32	10.3.1.1	0	100	0	i
*> 10.3.10.1/32	192.168.203.31	1000		0	64503 i
*> 10.4.1.1/32	0.0.0.0	0		32768	i
* 10.4.10.1/32	192.168.104.41	1001		0	64504 i
*>i	10.3.1.1	1000	100	0	64504 i
* i 10.5.100.1/32	172.16.5.1	0	1000	0	?
r>i 192.168.1.0	10.0.1.1	0	100	0	i
r>i 192.168.2.0	10.0.1.1	0	100	0	i
*>i 209.165.201.0/27	10.0.2.1	0	100	0	64497 64498 ?
*>i 209.165.202.128/27	10.0.2.1	0	100	0	64498 64495 64496 ?

Task 4: Configuring the BGP Community

Task Visual Objective

In this task, you will remove the BGP MED configuration that is enabled in the previous task. On the pod PE router, you will configure a service provider policy by using BGP communities. You will verify the service provider policy by sending BGP routes from the pod CE router with different BGP communities, and you will observe the BGP routing table on the pod PE routers.

The figure shows what you will accomplish in this task:



Activity Procedure

Complete these steps:

- Step 1** On the pod CE router, remove the route-map configuration that is applied to the BGP neighbors in the previous task.
- Step 2** The pod CE router should send BGP routes without BGP MED values. Verify that the pod PE router selects the direct path to the CE router.

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

```
<... output omitted ...>
```

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 10.3.1.1/32	0.0.0.0	0		32768	i
*> 10.3.10.1/32	192.168.103.31	0		0	64503 i
*> 10.4.1.1/32	10.4.1.1	0	100	0	i
*> 10.4.10.1/32	192.168.204.41	0		0	64504 i
* i10.5.100.1/32	172.16.5.1	0	1000	0	?
*> i192.168.1.0/24	10.0.1.1	0	100	0	i
*> i192.168.2.0/24	10.0.1.1	0	100	0	i
*> i209.165.201.0/27	10.0.2.1	0	100	0	64497 64498 ?
*> i209.165.202.128/27	10.0.2.1	0	100	0	64498 64495 64496 ?

```

PE4# show bgp
<... output omitted ...>

```

	Network	Next Hop	Metric	LocPrf	Weight	Path
r>i	10.3.1.1/32	10.3.1.1	0	100	0	i
* i	10.3.10.1/32	10.3.1.1	0	100	0	64503 i
*>		192.168.203.31	0		0	64503 i
*>	10.4.1.1/32	0.0.0.0	0		32768	i
* i	10.4.10.1/32	10.3.1.1	0	100	0	64504 i
*>		192.168.104.41	0		0	64504 i
* i	10.5.100.1/32	172.16.5.1	0	1000	0	?
r>i	192.168.1.0	10.0.1.1	0	100	0	i
r>i	192.168.2.0	10.0.1.1	0	100	0	i
*>i	209.165.201.0/27	10.0.2.1	0	100	0	64497 64498 ?
*>i	209.165.202.128/27	10.0.2.1	0	100	0	64498 64495 64496 ?

Step 3 On you team pod and the neighbor pod PE routers, use the RPL (Cisco IOS XR) or **route-map** command (Cisco IOS XE) to configure the BGP community policy, as shown in the table:

Community	Action
64500:101	Set local preference 101
64500:99	Set local preference 99
64500:999	Set MED 999
64500:1000	Set MED 1000
64500:1001	Set MED 1001
64500:1	Prepend AS-path one time
64500:2	Prepend AS-path two times
64500:3	Prepend AS-path three times

- Step 4** On the pod and neighbor pod PE routers, apply the RPL or route-map to the BGP neighbors (pod CE router).
- Step 5** Enable the pod CE router to send BGP communities and use the BGP community new format. Enable the pod PE (Cisco IOS XE only) router to send BGP communities toward the P1 router and to use the BGP community new format. The router running Cisco IOS XR is sending BGP communities by default.
- Step 6** On the pod CE router, configure route-maps Primary_SP and Backup_SP. The Primary_SP route-map will set BGP community 64500:101, and the Backup_SP route-map will set BGP community 64500:99. On the pod CE router, apply route-map Primary_SP to the primary BGP neighbor (neighbor pod PE router) and apply Backup_SP to the backup BGP neighbor (your pod PE router).
- Step 7** On the pod and neighbor pod PE routers, verify that BGP communities prepend to the BGP routes from the pod CE router. On the pod and neighbor pod PE routers, also verify the BGP routing table. Routes from the pod CE router have a different BGP local preference.

```
RP/0/RSP0/CPU0:PE3# show bgp 10.3.10.1/32
Mon Jun  9 13:32:04.096 UTC
BGP routing table entry for 10.3.10.1/32
Versions:
  Process          bRIB/RIB   SendTblVer
  Speaker          43         43
Last Modified: Jun  9 13:29:09.532 for 00:02:54
Paths: (2 available, best #1)
  Advertised to update-groups (with more than one peer):
    0.3
  Path #1: Received by speaker 0
  Advertised to update-groups (with more than one peer):
    0.3
  64503
    10.4.1.1 (metric 30) from 10.0.1.1 (10.4.1.1)
      Origin IGP, metric 0, localpref 101, valid, internal, best, group-best, import-
candidate, import suspect
      Received Path ID 0, Local Path ID 1, version 43
      Community: 64500:101
      Originator: 10.4.1.1, Cluster list: 10.0.1.1
  Path #2: Received by speaker 0
  Not advertised to any peer
  64503
    192.168.103.31 from 192.168.103.31 (10.3.10.1)
      Origin IGP, metric 0, localpref 99, valid, external, import suspect
      Received Path ID 0, Local Path ID 0, version 0
      Community: 64500:99
      Origin-AS validity: not-found
```

RP/0/RSP0/CPU0:PE3# show bgp

<... output omitted ...>

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 10.3.1.1/32	0.0.0.0	0		32768	i
*>i 10.3.10.1/32	10.4.1.1	0	101	0	64503 i
*	192.168.103.31	0	99	0	64503 i
*>i 10.4.1.1/32	10.4.1.1	0	100	0	i
*> 10.4.10.1/32	192.168.204.41	0	101	0	64504 i
* i 10.5.100.1/32	172.16.5.1	0	1000	0	?
*>i 192.168.1.0/24	10.0.1.1	0	100	0	i
*>i 192.168.2.0/24	10.0.1.1	0	100	0	i
*>i 209.165.201.0/27	10.0.2.1	0	100	0	64497 64498 ?
*>i 209.165.202.128/27	10.0.2.1	0	100	0	64498 64495 64496 ?

Processed 9 prefixes, 10 paths

PE4# show ip bgp 10.4.10.1/32

BGP routing table entry for 10.4.10.1/32, version 16

Paths: (2 available, best #1, table default)

Advertised to update-groups:

7

Refresh Epoch 1

64504

10.3.1.1 {metric 30} from 10.0.1.1 {10.0.1.1}

Origin IGP, metric 0, localpref 101, valid, internal, best

Community: 64500:101

Originator: 10.3.1.1, Cluster list: 10.0.1.1

rx pathid: 0, tx pathid: 0x0

Refresh Epoch 2

64504

192.168.104.41 from 192.168.104.41 {10.4.10.1}

Origin IGP, metric 0, localpref 99, valid, external

Community: 64500:99

rx pathid: 0, tx pathid: 0

PE4# show ip bgp

<... output omitted ...>

Network	Next Hop	Metric	LocPrf	Weight	Path
r>i 10.3.1.1/32	10.3.1.1	0	100	0	i
*> 10.3.10.1/32	192.168.203.31	0	101	0	64503 i
*> 10.4.1.1/32	0.0.0.0	0		32768	i
*>i 10.4.10.1/32	10.3.1.1	0	101	0	64504 i
*	192.168.104.41	0	99	0	64504 i
* i 10.5.100.1/32	172.16.5.1	0	1000	0	?
r>i 192.168.1.0	10.0.1.1	0	100	0	i
r>i 192.168.2.0	10.0.1.1	0	100	0	i
*>i 209.165.201.0/27	10.0.2.1	0	100	0	64497 64498 ?
*>i 209.165.202.128/27	10.0.2.1	0	100	0	64498 64495 64496 ?

- Step 8** On the pod CE router, modify the route-maps that are configured in the previous step. The primary_SP route-map should set BGP community 64500:999, and the Backup_SP route-map should set BGP community 64500:1000.
- Step 9** On the pod and the neighbor pod PE routers, verify the BGP routes from the pod CE router. Routes from the pod CE router have different BGP communities. On the pod and neighbor pod PE routers, verify the BGP routing table. Routes from the pod CE router have a different BGP MED (metric). On the pod CE router, verify the route that is received from the neighbor pod CE router. Observe the BGP AS-path attribute.

```
RP/0/RSP0/CPU0:PE3# show bgp 10.3.10.1/32
Mon Jun  9 13:40:24.299 UTC
BGP routing table entry for 10.3.10.1/32
Versions:
  Process          bRIB/RIB   SendTblVer
  Speaker          45         45
Last Modified: Jun  9 13:38:32.532 for 00:01:51
Paths: (2 available, best #1)
  Advertised to update-groups (with more than one peer):
    0.3
  Path #1: Received by speaker 0
  Advertised to update-groups (with more than one peer):
    0.3
  64503
    10.4.1.1 (metric 30) from 10.0.1.1 (10.4.1.1)
      Origin IGP, metric 999, localpref 100, valid, internal, best, group-best, import-candidate
      Received Path ID 0, Local Path ID 1, version 45
      Community: 64500:999
      Originator: 10.4.1.1, Cluster list: 10.0.1.1
  Path #2: Received by speaker 0
  Not advertised to any peer
  64503
    192.168.103.31 from 192.168.103.31 (10.3.10.1)
      Origin IGP, metric 1000, localpref 100, valid, external
      Received Path ID 0, Local Path ID 0, version 0
      Community: 64500:1000
      Origin-AS validity: not-found
```

```
RP/0/RSP0/CPU0:PE3# show bgp
<... output omitted ...>
      Network          Next Hop          Metric LocPrf Weight Path
*> 10.3.1.1/32         0.0.0.0              0          32768 i
*>i10.3.10.1/32        10.4.1.1              999         100      0 64503 i
*                      192.168.103.31        1000         0      0 64503 i
*>i10.4.1.1/32         10.4.1.1              0          100      0 i
*> 10.4.10.1/32        192.168.204.41        999         0      0 64504 i
* i10.5.100.1/32       172.16.5.1            0          1000     0 ?
*>i192.168.1.0/24      10.0.1.1              0          100      0 i
*>i192.168.2.0/24      10.0.1.1              0          100      0 i
*>i209.165.201.0/27    10.0.2.1              0          100      0 64497 64498 ?
*>i209.165.202.128/27 10.0.2.1              0          100      0 64498 64495 64496 ?

Processed 9 prefixes, 9 paths
```

```

PE4# show ip bgp 10.4.10.1/32
BGP routing table entry for 10.4.10.1/32, version 18
Paths: (2 available, best #1, table default)
  Advertised to update-groups:
    7
  Refresh Epoch 1
  64504
    10.3.1.1 (metric 30) from 10.0.1.1 (10.0.1.1)
      Origin IGP, metric 999, localpref 100, valid, internal, best
      Community: 64500:999
      Originator: 10.3.1.1, Cluster list: 10.0.1.1
      rx pathid: 0, tx pathid: 0x0
  Refresh Epoch 3
  64504
    192.168.104.41 from 192.168.104.41 (10.4.10.1)
      Origin IGP, metric 1000, localpref 100, valid, external
      Community: 64500:1000
      rx pathid: 0, tx pathid: 0

```

```

PE4# show ip bgp
<... output omitted ...>

```

Network	Next Hop	Metric	LocPrf	Weight	Path
r>i 10.3.1.1/32	10.3.1.1	0	100	0	i
*> 10.3.10.1/32	192.168.203.31	999		0	64503 i
*> 10.4.1.1/32	0.0.0.0	0		32768	i
*>i 10.4.10.1/32	10.3.1.1	999	100	0	64504 i
* 192.168.104.41	192.168.104.41	1000		0	64504 i
* i 10.5.100.1/32	172.16.5.1	0	1000	0	?
r>i 192.168.1.0	10.0.1.1	0	100	0	i
r>i 192.168.2.0	10.0.1.1	0	100	0	i
*>i 209.165.201.0/27	10.0.2.1	0	100	0	64497 64498 ?
*>i 209.165.202.128/27	10.0.2.1	0	100	0	64498 64495 64496 ?

```

CE3# show ip bgp
<... output omitted ...>

```

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 10.3.1.1/32	192.168.203.40			2	64500 i
* 192.168.103.30	192.168.103.30	0		1	64500 i
*> 10.3.10.1/32	0.0.0.0	0		32768	i
* 10.4.1.1/32	192.168.103.30			1	64500 i
*> 192.168.203.40	192.168.203.40	0		2	64500 i
* 10.4.10.1/32	192.168.103.30			1	64500 64504 i
*> 192.168.203.40	192.168.203.40			2	64500 64504 i
* 192.168.1.0	192.168.103.30			1	64500 i
*> 192.168.203.40	192.168.203.40			2	64500 i
* 192.168.2.0	192.168.103.30			1	64500 i
*> 192.168.203.40	192.168.203.40			2	64500 i
* 209.165.201.0/27	192.168.103.30			1	64500 64497 64498 ?
*> 192.168.203.40	192.168.203.40			2	64500 64497 64498 ?
* 209.165.202.128/27	192.168.103.30			1	64500 64498 64495 64496
? 192.168.203.40	192.168.203.40			2	64500 64498 64495 64496
? 192.168.203.40	192.168.203.40			2	64500 64498 64495 64496

```

CE4# show ip bgp
<... output omitted ...>

```

Network	Next Hop	Metric	LocPrf	Weight	Path
* 10.3.1.1/32	192.168.104.40			1	64500 i
*> 10.3.1.1/32	192.168.204.30	0		2	64500 i
*> 10.3.10.1/32	192.168.204.30			2	64500 64503 i
* 10.3.10.1/32	192.168.104.40			1	64500 64503 i
*> 10.4.1.1/32	192.168.204.30			2	64500 i
* 10.4.1.1/32	192.168.104.40	0		1	64500 i
*> 10.4.10.1/32	0.0.0.0	0		32768	i
*> 192.168.1.0	192.168.204.30			2	64500 i
* 192.168.1.0	192.168.104.40			1	64500 i
*> 192.168.2.0	192.168.204.30			2	64500 i
* 192.168.2.0	192.168.104.40			1	64500 i
*> 209.165.201.0/27	192.168.204.30			2	64500 64497 64498 ?
* 209.165.201.0/27	192.168.104.40			1	64500 64497 64498 ?
*> 209.165.202.128/27	192.168.204.30			2	64500 64498 64495 64496
? 209.165.202.128/27	192.168.104.40			1	64500 64498 64495 64496
? 209.165.202.128/27					

Step 10 On the pod CE router, modify the route-maps that are configured in the previous step. The Primary_SP route-map should set BGP community 64500:1, and the Backup_SP route-map should set BGP community 64500:3.

When using AS-path prepending, you are trying to influence the return traffic from your neighbor AS. In this case, you can see from your neighbor CE routing table that they have two routes to your CE router loopback interface. For example, from the CE3 **show ip bgp** command output, there are two paths to the CE4 loopback interface. The best path that is selected is the one with the higher weight, although it has a longer AS path. Using AS-path prepending to try to influence the return path from your neighbor CE router did not work in this case because weight is examined first before the other attributes in determining the best path to use.

Verification

You have completed this task when you attain this result:

- Step 1** On the pod and neighbor pod PE routers, verify the BGP routes from the pod CE router. Routes from the pod CE router have different BGP communities. On the pod and neighbor pod PE routers, verify the BGP routing table. Observe routes from the pod CE router. On the pod CE router, verify the routes that are received from the neighbor pod CE router. Observe that the BGP AS-path attribute is different for two routes.

```
RP/0/RSP0/CPU0:PE3# show bgp 10.3.10.1/32
Mon Jun  9 13:53:32.486 UTC
BGP routing table entry for 10.3.10.1/32
Versions:
  Process          bRIB/RIB  SendTblVer
  Speaker          47        47
Last Modified: Jun  9 13:47:25.532 for 00:06:07
Paths: (1 available, best #1)
  Advertised to update-groups (with more than one peer):
    0.3
  Advertised to peers (in unique update groups):
    10.0.1.1
  Path #1: Received by speaker 0
  Advertised to update-groups (with more than one peer):
    0.3
  Advertised to peers (in unique update groups):
    10.0.1.1
  64503
    192.168.103.31 from 192.168.103.31 (10.3.10.1)
      Origin IGP, metric 0, localpref 100, valid, external, best, group-best, import-
candidate, import suspect
      Received Path ID 0, Local Path ID 1, version 47
      Community: 64500:3
      Origin-AS validity: not-found
```

```
RP/0/RSP0/CPU0:PE3# show bgp
<... output omitted ...>
```

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 10.3.1.1/32	0.0.0.0	0		32768	i
*> 10.3.10.1/32	192.168.103.31	0		0	64503 i
*>i10.4.1.1/32	10.4.1.1	0	100	0	i
*> 10.4.10.1/32	192.168.204.41	0		0	64504 i
* i10.5.100.1/32	172.16.5.1	0	1000	0	?
*>i192.168.1.0/24	10.0.1.1	0	100	0	i
*>i192.168.2.0/24	10.0.1.1	0	100	0	i
*>i209.165.201.0/27	10.0.2.1	0	100	0	64497 64498 ?
*>i209.165.202.128/27	10.0.2.1	0	100	0	64498 64495 64496 ?

Processed 9 prefixes, 9 paths

```

PE4# show ip bgp 10.4.10.1/32
BGP routing table entry for 10.4.10.1/32, version 20
Paths: (2 available, best #2, table default)
  Advertised to update-groups:
    7      8
  Refresh Epoch 1
64504
  10.3.1.1 (metric 30) from 10.0.1.1 (10.0.1.1)
    Origin IGP, metric 0, localpref 100, valid, internal
    Community: 64500:1
    Originator: 10.3.1.1, Cluster list: 10.0.1.1
    rx pathid: 0, tx pathid: 0
  Refresh Epoch 4
64504
  192.168.104.41 from 192.168.104.41 (10.4.10.1)
    Origin IGP, metric 0, localpref 100, valid, external, best
    Community: 64500:3
    rx pathid: 0, tx pathid: 0x0

```

```

PE4# show ip bgp
<... output omitted ...>

```

Network	Next Hop	Metric	LocPrf	Weight	Path
r>i 10.3.1.1/32	10.3.1.1	0	100	0	i
* i 10.3.10.1/32	10.3.1.1	0	100	0	64503 i
*>	192.168.203.31	0		0	64503 i
*> 10.4.1.1/32	0.0.0.0	0		32768	i
* i 10.4.10.1/32	10.3.1.1	0	100	0	64504 i
*>	192.168.104.41	0		0	64504 i
* i 10.5.100.1/32	172.16.5.1	0	1000	0	?
r>i 192.168.1.0	10.0.1.1	0	100	0	i
r>i 192.168.2.0	10.0.1.1	0	100	0	i
*>i 209.165.201.0/27	10.0.2.1	0	100	0	64497 64498 ?
*>i 209.165.202.128/27	10.0.2.1	0	100	0	64498 64495 64496 ?

```

CE3# show ip bgp
<... output omitted ...>

```

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 10.3.1.1/32	192.168.203.40			2	64500 i
*	192.168.103.30	0		1	64500 i
*> 10.3.10.1/32	0.0.0.0	0		32768	i
*	192.168.103.30			1	64500 i
*>	192.168.203.40	0		2	64500 i
*	10.4.10.1/32			1	64500 64500 64504 i
*>	192.168.203.40			2	64500 64500 64500 64500
64504 i					
*	192.168.1.0			1	64500 i
*>	192.168.203.40			2	64500 i
*	192.168.2.0			1	64500 i
*>	192.168.203.40			2	64500 i
*	209.165.201.0/27			1	64500 64497 64498 ?
*>	192.168.203.40			2	64500 64497 64498 ?
*	209.165.202.128/27				
	192.168.103.30			1	64500 64498 64495 64496
?					
*>	192.168.203.40			2	64500 64498 64495 64496
?					

```

CE4# show ip bgp
<... output omitted ...>
  Network      Next Hop      Metric LocPrf Weight Path
  * 10.3.1.1/32 192.168.104.40          0      1 64500 i
  *>          192.168.204.30          0      2 64500 i
  *> 10.3.10.1/32 192.168.204.30          0      2 64500 64500 64500 64500
64503 i
  *          192.168.104.40          0      1 64500 64500 64503 i
  *> 10.4.1.1/32 192.168.204.30          0      2 64500 i
  *          192.168.104.40          0      1 64500 i
  *> 10.4.10.1/32 0.0.0.0          0 32768 i
  *> 209.165.201.0/27 192.168.204.30          0      2 64500 64497 64498 ?
  *          192.168.104.40          0      1 64500 64497 64498 ?
  *> 209.165.202.128/27 192.168.204.30          0      2 64500 64498 64495 64496
?

  * 10.3.1.1/32 192.168.104.40          0      1 64500 i
  *>          192.168.204.30          0      2 64500 i
  *> 10.3.10.1/32 192.168.204.30          0      2 64500 64500 64500 64500
64503 i
  *          192.168.104.40          0      1 64500 64500 64503 i
  *> 10.4.1.1/32 192.168.204.30          0      2 64500 i
  *          192.168.104.40          0      1 64500 i
  *> 10.4.10.1/32 0.0.0.0          0 32768 i
  *> 192.168.1.0 192.168.204.30          0      2 64500 i
  *          192.168.104.40          0      1 64500 i
  *> 192.168.2.0 192.168.204.30          0      2 64500 i
  *          192.168.104.40          0      1 64500 i
  *> 209.165.201.0/27 192.168.204.30          0      2 64500 64497 64498 ?
  *          Network      Next Hop      Metric LocPrf Weight Path
  *          192.168.104.40          0      1 64500 64497 64498 ?
  *> 209.165.202.128/27 192.168.204.30          0      2 64500 64498 64495 64496
?
  *          192.168.104.40          0      1 64500 64498 64495 64496
?

  *          192.168.104.40          0      1 64500 64498 64495 64496
?

```


Answer Key

Hardware Lab 1: Implement OSPF Routing

Task 1: Enabling OSPF on the Routers

Step 1 Enter these commands on the CE and PE routers, running IOS, IOS XR, and IOS XE, respectively.

```
CEx(config)#
interface Loopback0
 ip address 10.x.10.1 255.255.255.255
 no shutdown
!
interface GigabitEthernet0/0
 ip address 192.168.10x.x1 255.255.255.0
 no shutdown
```

```
RP/0/RSP0/CPU0: PEx(config)#
interface Loopback0
 ipv4 address 10.x.1.1 255.255.255.255
 no shutdown
!
interface GigabitEthernet0/0/0/0
 ipv4 address 192.168.10x.x0 255.255.255.0
 no shutdown
!
interface GigabitEthernet0/0/0/1
 ipv4 address 192.168.lxy.x0 255.255.255.0
 no shutdown
!
interface GigabitEthernet0/0/0/2
 ipv4 address 192.168.x1.x0 255.255.255.0
 no shutdown
!
interface GigabitEthernet0/0/0/3
 ipv4 address 192.168.x2.x0 255.255.255.0
 no shutdown
!
commit
```

```

CEy(config)#
interface Loopback0
 ip address 10.y.10.1 255.255.255.255
 no shutdown
!
interface GigabitEthernet0/0
 ip address 192.168.10y.y1 255.255.255.0
 no shutdown

```

```

PEy(config)#
interface Loopback0
 ip address 10.y.1.1 255.255.255.255
 no shutdown
!
interface GigabitEthernet0/0/0
 ip address 192.168.10y.y0 255.255.255.0
 no shutdown
!
interface GigabitEthernet0/0/1
 ip address 192.168.1xy.y0 255.255.255.0
 no shutdown
!
interface GigabitEthernet0/0/2
 ip address 192.168.y1.y0 255.255.255.0
 no shutdown
!
interface GigabitEthernet0/0/3
 ip address 192.168.y2.y0 255.255.255.0
 no shutdown

```

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

Step 3 The following table shows the OSPF areas that are assigned to the interfaces in your team.

Link (Interface)	OSPF Area
PE _x (Gi0/0/0/2) – P1	0
PE _x (Gi0/0/0/3) – P2	0
PE _x (Gi0/0/0/1) – PE _y (Gi0/0/0/1)	0
PE _x (Gi0/0/0/0) – CE _x (Gi0/0)	1 _x
PE _x (Loopback0)	x
CE _x (Loopback0)	1 _x
PE _y (Gi0/0/0/2) – P1	0
PE _y (Gi0/0/0/3) – P2	0
PE _y (Gi0/0/0/0) – CE _y (Gi0/0)	1 _y

Link (Interface)	OSPF Area
PEy (Loopback0)	y
CEy (Loopback0)	1y

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

Step 4 Enter these commands on the CE and PE routers, running IOS, IOS XR, and IOS XE, respectively.

```
CEx(config)#
router ospf 1
 network 10.x.10.1 0.0.0.0 area 1x
 network 192.168.10x.x1 0.0.0.0 area 1x
```

```
RP/0/RSP0/CPU0: PEx(config)#
router ospf 1
 address-family ipv4
 area 0
 interface GigabitEthernet0/0/0/1
 interface GigabitEthernet0/0/0/2
 interface GigabitEthernet0/0/0/3
 area x
 interface Loopback0
 area 1x
 interface GigabitEthernet0/0/0/0
 !
commit
```

```
CEy(config)#
router ospf 1
 network 10.y.10.1 0.0.0.0 area 1y
 network 192.168.10y.y1 0.0.0.0 area 1y
```

```
PEy(config)#
router ospf 1
 network 10.y.1.1 0.0.0.0 area y
 network 192.168.y1.y0 0.0.0.0 area 0
 network 192.168.y2.y0 0.0.0.0 area 0
 network 192.168.10y.y0 0.0.0.0 area 1y
 network 192.168.1xy.y0 0.0.0.0 area 0
```

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

Step 6 The following table shows the router ID that is assigned to the OSPF routing processes in a team.

Router	OSPF Router ID
PE _x	10.x.1.1
CE _x	10.x.10.1
PE _y	10.y.1.1
CE _y	10.y.10.1

x should be substituted with 1 or 3, and *y* should be substituted with 2 or 4, depending on your team that is assigned.

Step 7 Enter these commands on the CE and PE routers, running IOS, IOS XR, and IOS XE, respectively.

```
CEx(config)#  
router ospf 1  
  router-id x0.x0.x0.x0  
end  
!  
clear ip ospf process  
Reset ALL OSPF processes? [no]: yes
```

```
RP/0/RSP0/CPU0:PEx(config)#  
router ospf 1  
  router-id x.x.x.x  
!  
commit  
!  
end  
!  
clear ospf process  
Reset ALL OSPF processes? [no]: yes
```

```
CEy(config)#  
router ospf 1  
  router-id y0.y0.y0.y0  
end  
!  
clear ip ospf process  
Reset ALL OSPF processes? [no]: yes
```

```
PEy(config)#  
router ospf 1  
  router-id y.y.y.y  
end  
!  
clear ip ospf process  
Reset ALL OSPF processes? [no]: yes
```

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

Step 8 Enter these commands on the CE and PE routers, running IOS, IOS XR, and IOS XE, respectively.

```
CEx(config)#
ipv6 unicast-routing
!
interface Loopback0
  ipv6 address 2001:DB8:10:x:10::1/128
  ipv6 enable
!
interface GigabitEthernet0/0
  ipv6 enable
```

```
RP/0/RSP0/CPU0:PEx(config)#
interface Loopback0
  ipv6 address 2001:db8:10:x:1::1/128
  ipv6 enable
!
interface GigabitEthernet0/0/0/0
  ipv6 enable
!
interface GigabitEthernet0/0/0/1
  ipv6 enable
!
interface GigabitEthernet0/0/0/2
  ipv6 enable
!
interface GigabitEthernet0/0/0/3
  ipv6 enable
!
commit
```

```
CEy(config)#
ipv6 unicast-routing
!
interface Loopback0
  ipv6 address 2001:DB8:10:y:10::1/128
  ipv6 enable
!
interface GigabitEthernet0/0
  ipv6 enable
```

```

PEy(config)#
ipv6 unicast-routing
!
interface Loopback0
  ipv6 address 2001:DB8:10:y:1::1/128
  ipv6 enable
!
interface GigabitEthernet0/0/0
  ipv6 enable
!
interface GigabitEthernet0/0/1
  ipv6 enable
!
interface GigabitEthernet0/0/2
  ipv6 enable
!
interface GigabitEthernet0/0/3
  ipv6 enable

```

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

Step 9 Enter these commands on the CE and PE routers, running IOS, IOS XR, and IOS XE, respectively.

```

CEx(config)#
interface Loopback0
  ipv6 ospf 1 area 1x
!
interface GigabitEthernet0/0
  ipv6 ospf 1 area 1x

```

```

RP/0/RSP0/CPU0:PEx(config)#
router ospfv3 1
  area 0
    interface GigabitEthernet0/0/0/1
    !
    interface GigabitEthernet0/0/0/2
    !
    interface GigabitEthernet0/0/0/3
    !
  area x
    interface Loopback0
    !
  area 1x
    interface GigabitEthernet0/0/0/0
    !
commit

```

```

CEy(config)#
interface Loopback0
  ipv6 ospf 1 area 1y
!
interface GigabitEthernet0/0
  ipv6 ospf 1 area 1y

```

```

PEy(config)#
interface Loopback0
  ipv6 ospf 1 area y
!
interface GigabitEthernet0/0/0
  ipv6 ospf 1 area 1y
!
interface GigabitEthernet0/0/1
  ipv6 ospf 1 area 0
!
interface GigabitEthernet0/0/2
  ipv6 ospf 1 area 0
!
interface GigabitEthernet0/0/3
  ipv6 ospf 1 area 0

```

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

Task 2: Influencing OSPF DR and BDR Election

Step 1 Enter these commands on the CE routers:

```

CEx(config)#
interface GigabitEthernet0/0
  ip ospf priority 0

```

```

CEy(config)#
interface GigabitEthernet0/0
  ip ospf priority 0

```

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

Step 2 Enter these commands on the PE routers, running IOS XR and IOS XE, respectively.

```

RP/0/RSP0/CPU0: PEx(config)#
router ospf 1
 area 0
   interface GigabitEthernet0/0/0/2
     priority 2
   interface GigabitEthernet0/0/0/3
     priority 2
!
commit

```

```

PEy(config)#
interface GigabitEthernet0/0/2
  ip ospf priority 2
interface GigabitEthernet0/0/3
  ip ospf priority 2

```

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

Task 3: Influencing OSPF Route Selection by Changing the OSPF Link Cost

Step 2 Enter these commands on the PE routers, running IOS XR and IOS XE, respectively.

```
RP/0/RSP0/CPU0:PEx(config)#  
router ospf 1  
 area 0  
   interface GigabitEthernet0/0/0/3  
     cost 3  
!  
commit
```

```
PEy(config)#  
interface GigabitEthernet0/0/2  
 ip ospf cost 3
```

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

Step 3 Enter these commands on the PE routers, running IOS XR and IOS XE, respectively.

```
RP/0/RSP0/CPU0:PEx(config)#  
router ospf 1  
 area 0  
   interface GigabitEthernet0/0/0/1  
     cost 10  
!  
commit
```

```
PEy(config)#  
interface GigabitEthernet0/0/1  
 ip ospf cost 10
```

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

Task 4: Configuring OSPF Authentication

Step 1 Enter these commands on the CE and PE routers, running IOS, IOS XR, and IOS XE, respectively.

```
CEx(config)#  
interface GigabitEthernet0/0  
 ip ospf authentication message-digest  
 ip ospf message-digest-key 1 md5 cisco
```

```
RP/0/RSP0/CPU0:PEx(config)#
router ospf 1
 area 1x
   interface GigabitEthernet0/0/0/0
     authentication message-digest
     message-digest-key 1 md5 cisco
!
commit
```

```
CEy(config)#
interface GigabitEthernet0/0
 ip ospf authentication message-digest
 ip ospf message-digest-key 1 md5 cisco
```

```
PEy(config)#
interface GigabitEthernet0/0/0
 ip ospf authentication message-digest
 ip ospf message-digest-key 1 md5 cisco
```

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

Task 5: Configuring OSPF Virtual Links

Step 1 Enter these commands on the CE routers.

```
CEx(config)#
interface Loopback1x
 ip address 1x.0.0.1 255.255.255.0
```

```
CEy(config)#
interface Loopback1y
 ip address 1y.0.0.1 255.255.255.0
```

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

Step 2 Enter these commands on the CE routers.

```
CEx(config)#
router ospf 1
 network 1x.0.0.1 0.0.0.0 area x00
```

```
CEy(config)#
router ospf 1
 network 1y.0.0.1 0.0.0.0 area y00
```

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

Step 3 Enter these commands on the CE and PE routers, running IOS, IOS XR, and IOS XE, respectively.

```
CEx(config)#  
router ospf 1  
area 1x virtual-link x.x.x.x
```

```
RP/0/RSP0/CPU0:PEx(config)#  
router ospf 1  
area 1x  
virtual-link x0.x0.x0.x0  
!  
commit
```

```
CEy(config)#  
router ospf 1  
area 1y virtual-link y.y.y.y
```

```
PEy(config)#  
router ospf 1  
area 1y virtual-link y0.y0.y0.y0
```

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

Hardware Lab 2: Implement OSPF Special Area Types

Task 1: Enabling OSPF Summarization

Step 1 Remove the OSPF virtual link on the CE and PE routers:

```
CEx(config)#  
router ospf 1  
no area 1x virtual-link x.x.x.x  
no network 1x.0.0.1 0.0.0.0 area x00
```

```
RP/0/RSP0/CPU0:PEx(config)#  
router ospf 1  
area 1x  
no virtual-link x0.x0.x0.x0  
!  
commit
```

```
CEy(config)#  
router ospf 1  
no area 1y virtual-link y.y.y.y  
no network 1y.0.0.1 0.0.0.0 area y00
```

```
PEy(config)#
router ospf 1
no area 1y virtual-link y0.y0.y0.y0
```

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

Step 2 Redistribute Loopback1x or Loopback1y into the OSPF process on the CE router:

```
CEx(config)#
router ospf 1
redistribute connected subnets route-map Loopback1x
!
route-map Loopback1x permit 10
match interface Loopback1x
```

```
CEy(config)#
router ospf 1
redistribute connected subnets route-map Loopback1y
!
route-map Loopback1y permit 10
match interface Loopback1y
```

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

Step 3 Configure OSPF summarization on the PE router:

```
RP/0/RSP0/CPU0:PEx(config)#
router ospf 1
area 1x
range 10.x.10.0/24
!
commit
```

```
PEy(config)#
router ospf 1
area 1y range 10.y.10.0 255.255.255.0
```

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

Step 4 Configure OSPF summarization on the CE router:

```
CEx(config)#
router ospf 1
summary-address 1x.0.0.0 255.0.0.0
```

```
CEy(config)#
router ospf 1
summary-address 1y.0.0.0 255.0.0.0
```

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

Task 2: Configuring the OSPF Stub Area

Step 2 Configure the OSPF stub area on the CE and PE routers:

```
CEx(config)#  
router ospf 1  
area 1x stub
```

```
RP/0/RSP0/CPU0:PEx(config)#  
router ospf 1  
area 1x  
stub  
!  
commit
```

```
CEy(config)#  
router ospf 1  
area 1y stub
```

```
PEy(config)#  
router ospf 1  
area 1y stub
```

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

Step 4 Configure the OSPF totally stubby area on the PE router:

```
RP/0/RSP0/CPU0:PEx(config)#  
router ospf 1  
area 1x  
stub no-summary  
!  
commit
```

```
PEy(config)#  
router ospf 1  
area 1y stub no-summary
```

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

Step 6 Adjust the OSPF cost of the default route on the PE router:

```
RP/0/RSP0/CPU0:PEx(config)#
router ospf 1
 area 1x
  default-cost 10
!
commit
```

```
PEy(config)#
router ospf 1
 area 1y default-cost 10
```

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

Step 9 Configure the totally NSSA on the CE and PE routers:

```
CEx(config)#
router ospf 1
 no area 1x stub
 area 1x nssa
```

```
RP/0/RSP0/CPU0:PEx(config)#
router ospf 1
 area 1x
  no stub
  nssa no-summary
!
commit
```

```
CEy(config)#
router ospf 1
 no area 1y stub
 area 1y nssa
```

```
PEy(config)#
router ospf 1
 no area 1y stub
 area 1y nssa no-summary
```

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

Hardware Lab 3: Implement Integrated IS-IS Routing

Task 1: Enabling Integrated IS-IS on the Routers

Step 3 The following table shows IS-IS NET addressing and interfaces that are used in IS-IS routing.

Router	IS-IS Area	System ID+00	IS-IS Interfaces
CE (your pod)	49.000x	0.100.0x01.0001.00	Gi0/0, Lo0
PE (your pod)	49.000x	0.100.0x00.0001.00	Gi0/0/0/0, Gi0/0/0/1, Lo0
CE (neighbor pod)	49.000y	0.100.0y01.0001.00	Gi0/0, Lo0
PE (neighbor pod)	49.000y	0.100.0y00.0001.00	Gi0/0/0, Gi0/0/1, Lo0

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

Step 4 The Integrated IS-IS configuration is on the CE and PE routers:

```
CEx(config)#
no router ospf 1
!
interface Loopback0
ip router isis
!
interface GigabitEthernet0/0
ip router isis
!
router isis
net 49.000x.0100.0x01.0001.00
```

```
RP/0/RSP0/CPU0:PEx(config)#
no router ospf 1
!
router isis 1
net 49.000x.0100.0x00.1001.00
!
interface Loopback0
address-family ipv4 unicast
!
interface GigabitEthernet0/0/0/0
address-family ipv4 unicast
!
interface GigabitEthernet0/0/0/1
address-family ipv4 unicast
!
commit
```

```

CEy(config)#
no router ospf 1
!
interface Loopback0
 ip router isis
!
interface GigabitEthernet0/0
 ip router isis
!
router isis
 net 49.000y.0100.0y01.0001.00

```

```

PEy(config)#
no router ospf 1
!
interface Loopback0
 ip router isis
!
interface GigabitEthernet0/0/0
 ip router isis
!
interface GigabitEthernet0/0/1
 ip router isis
!
router isis
 net 49.000y.0100.0y00.1001.00

```

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

Step 5 The IS-IS adjustments are configured on the CE and PE routers:

```

CEx(config)#
router isis
 is-type level-1

```

```

RP/0/RSP0/CPU0:PEx(config)#
router isis 1
 interface GigabitEthernet0/0/0/0
  circuit-type level-1
!
 interface GigabitEthernet0/0/0/1
  circuit-type level-2-only
!
commit

```

```

CEy(config)#
router isis
 is-type level-1

```

```

PEy(config)#
interface GigabitEthernet0/0/0
isis circuit-type level-1
!
interface GigabitEthernet0/0/1
isis circuit-type level-2-only

```

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

Task 2: IS-IS Route Summarization

Step 1 The IPv6 IS-IS configuration is on the CE and PE routers:

```

RP/0/RSP0/CPU0:PEx(config)#
router isis 1
address-family ipv4 unicast
summary-prefix 10.x.10.0/24 level 2
!
commit

```

```

PEy(config)#
router isis
summary-address 10.y.10.0 255.255.255.0

```

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

Task 3: Enabling IPv6 IS-IS Routing

Step 2 The IPv6 IS-IS configuration is on the CE and PE routers:

```

CEx(config)#
no ipv6 router ospf 1
!
interface Loopback0
ipv6 router isis
!
interface GigabitEthernet0/0
ipv6 router isis

```

```

RP/0/RSP0/CPU0:PEx(config)#
no router ospfv3 1
!
router isis 1
 address-family ipv6 unicast
  single-topology
!
interface Loopback0
 address-family ipv6 unicast
!
interface GigabitEthernet0/0/0/0
 address-family ipv6 unicast
!
interface GigabitEthernet0/0/0/1
 address-family ipv6 unicast
!
commit

```

```

CEy(config)#
no ipv6 router ospf 1
!
interface Loopback0
 ipv6 router isis
!
interface GigabitEthernet0/0
 ipv6 router isis

```

```

PEy(config)#
no ipv6 router ospf 1
!
interface Loopback0
 ipv6 router isis
!
interface GigabitEthernet0/0/0
 ipv6 router isis
!
interface GigabitEthernet0/0/1
 ipv6 router isis

```

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

Hardware Lab 4: Implement Basic BGP Routing

Task 1: Configuring External BGP

Step 1 The BGP configuration is on the PE routers:

```
RP/0/RSP0/CPU0:PEx(config)#  
router bgp 64500  
  address-family ipv4 unicast  
  !  
  neighbor 192.168.10x.x1  
    remote-as 6450x  
  address-family ipv4 unicast  
  !  
commit
```

```
PEy(config)#  
router bgp 64500  
  neighbor 192.168.10y.y1 remote-as 6450y
```

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

Step 2 The BGP configuration is on the CE routers:

```
CEx(config)#  
router bgp 6450x  
  neighbor 192.168.10x.x0 remote-as 64500
```

```
CEy(config)#  
router bgp 6450y  
  neighbor 192.168.10y.y0 remote-as 64500
```

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

Step 4 BGP authentication is on the CE and PE routers:

```
CEx(config)#  
router bgp 6450x  
  neighbor 192.168.10x.x0 password cisco
```

```
RP/0/RSP0/CPU0:PEx(config)#  
router bgp 64500  
  neighbor 192.168.10x.x1  
    password clear cisco  
  !  
commit
```

```
CEy(config)#
router bgp 6450y
neighbor 192.168.10y.y0 password cisco
```

```
PEy(config)#
router bgp 64500
neighbor 192.168.10y.y1 password cisco
```

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

Step 5 Advertise Loopback0 into BGP on the CE and PE routers:

```
CEx(config)#
router bgp 6450x
network 10.x.10.1 mask 255.255.255.255
```

```
RP/0/RSP0/CPU0: PEx(config)#
router bgp 64500
address-family ipv4 unicast
network 10.x.1.1/32
!
commit
```

```
CEy(config)#
router bgp 6450y
network 10.y.10.1 mask 255.255.255.255
```

```
PEy(config)#
router bgp 64500
network 10.y.1.1 mask 255.255.255.255
```

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

Step 6 Configure and apply inbound and outbound BGP route policy on the PE router.

```
RP/0/RSP0/CPU0: PEx(config)#
route-policy pass
pass
end-policy
!
router bgp 64500
neighbor 192.168.10x.x1
address-family ipv4 unicast
route-policy pass in
route-policy pass out
!
commit
```

x should be substituted with 1 or 3, depending on your team that is assigned.

Task 2: Configuring Internal BGP

Step 1 On the PE router, enable the interface and IS-IS routing toward the P1 router:

```
RP/0/RSP0/CPU0:PE3(config)#  
interface GigabitEthernet0/0/0/2  
  cdp  
  ipv4 address 192.168.x1.x0 255.255.255.0  
  ipv6 enable  
  no shutdown  
!  
router isis 1  
  interface GigabitEthernet0/0/0/2  
    circuit-type level-2-only  
    address-family ipv4 unicast  
  !  
  address-family ipv6 unicast  
!  
commit
```

```
PEy(config)#  
interface GigabitEthernet0/0/2  
  ip address 192.168.y1.y0 255.255.255.0  
  ip router isis  
  ipv6 enable  
  ipv6 router isis  
  cdp enable  
  isis circuit-type level-2-only
```

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

On the PE router, you should see the P1 router as an IS-IS neighbor.

```
RP/0/RSP0/CPU0:PE3# show isis neighbors  
Tue May 27 13:14:22.427 UTC
```

IS-IS 1 neighbors:

System Id	Interface	SNPA	State	Holdtime	Type	IETF-NSF
CE3	Gi0/0/0/0	2c54.2d9b.d3e0	Up	23	L1	Capable
P1	Gi0/0/0/2	6c9c.ed24.19ab	Up	20	L2	Capable
PE4	Gi0/0/0/1	442b.03e7.9381	Up	22	L2	Capable

Total neighbor count: 3

```
PE4# show isis neighbors
```

Tag null:

System Id	Type	Interface	IP Address	State	Holdtime	Circuit Id
CE4	L1	Gi0/0/0	192.168.104.41	UP	8	CE4.01
P1	L2	Gi0/0/2	192.168.41.1	UP	7	P1.01
PE3	L2	Gi0/0/1	192.168.134.30	UP	8	PE3.01

Step 2 The BGP configuration is on the PE routers:

```
RP/0/RSP0/CPU0:PEx(config)#
router bgp 64500
 neighbor 10.0.1.1
   remote-as 64500
   update-source Loopback0
 address-family ipv4 unicast
!
commit
```

```
PEy(config)#
router bgp 64500
 neighbor 10.0.1.1 remote-as 64500
 neighbor 10.0.1.1 update-source Loopback0
```

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

Step 3 The BGP next-hop self is on the PE router:

```
RP/0/RSP0/CPU0:PEx(config)#
router bgp 64500
 neighbor 10.0.1.1
   address-family ipv4 unicast
   next-hop-self
!
commit
```

```
PEy(config)#
router bgp 64500
 neighbor 10.0.1.1 next-hop-self
```

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

Before continuing to the verification section, make sure that the PE router in the neighbor pod (same team) finished the previous step.

Hardware Lab 5: Implement Route Redistribution

Task 1: Redistributing between OSPF and IS-IS

Step 1 Disable IS-IS and start OSPF on the CE routers:

```
CEx(config)#
no router isis
router ospf 1
 network 192.168.10x.x1 0.0.0.0 area 1x
 network 10.x.10.1 0.0.0.0 area 1x
```

```
CEy(config)#
no router isis
router ospf 1
 network 192.168.10y.y1 0.0.0.0 area 1y
 network 10.y.10.1 0.0.0.0 area 1y
```

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

Step 2 Disable IS-IS and start OSPF on the PE routers:

```
RP/0/RSP0/CPU0:PEx(config)#
router isis 1
 no interface GigabitEthernet0/0/0/0
 no interface GigabitEthernet0/0/0/1
!
router ospf 1
 area 1x
 interface GigabitEthernet0/0/0/0
 authentication message-digest
 message-digest-key 1 md5 cisco
 area 0
 interface GigabitEthernet0/0/0/1
!
commit
```

```
PEy(config)#
interface GigabitEthernet0/0/0
 no ip router isis
 no isis circuit-type level-1
!
interface GigabitEthernet0/0/1
 no ip router isis
 no isis circuit-type level-2-only
!
router ospf 1
 network 192.168.10y.y0 0.0.0.0 area 1y
 network 192.168.1xy.y0 0.0.0.0 area 0
```

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

Step 4 Redistribute between IS-IS and OSPF on the pod PE routers:

```
RP/0/RSP0/CPU0:PEx(config)#
router isis 1
  address-family ipv4 unicast
    redistribute ospf 1 level-2
  !
router ospf 1
  address-family ipv4 unicast
    redistribute isis 1 level-2
  !
commit
```

```
PEy(config)#
router ospf 1
  redistribute isis level-2 subnets
  !
router isis
  redistribute ospf 1 level-2
```

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

Task 2: Preventing Potential Routing Loops

Step 1 Configure the RPL or route map on the PE routers:

```
RP/0/RSP0/CPU0:PEx(config)#
route-policy OSPFintoISIS
  if tag eq 115 then drop else set tag 110 endif
exit
!
commit
```

```
PEy(config)#
route-map OSPFintoISIS deny 10
  match tag 115
  !
route-map OSPFintoISIS permit 20
  set tag 110
```

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

Step 2 Apply the RPL or route map to the OSPF into IS-IS redistribution:

```
RP/0/RSP0/CPU0:PEx(config)#
router isis 1
  address-family ipv4 unicast
    no redistribute ospf 1 level-2
    redistribute ospf 1 level-2 route-policy OSPFintoISIS
  !
commit
```

```
PEy(config)#
router isis
no redistribute ospf 1
redistribute ospf 1 route-map OSPFintoISIS
```

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

Step 3 Configure the RPL or route map on the PE routers:

```
RP/0/RSP0/CPU0:PEx(config)#
route-policy ISISintoOSPF
if tag eq 110 then drop else set tag 115 endif
exit
!
commit
```

```
PEy(config)#
route-map ISISintoOSPF deny 10
match tag 110
!
route-map ISISintoOSPF permit 20
set tag 115
```

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

Step 4 Apply the RPL or route map to the IS-IS into OSPF redistribution:

```
RP/0/RSP0/CPU0:PEx(config)#
router ospf 1
no redistribute isis 1 level-2
redistribute isis 1 level-2 route-policy ISISintoOSPF
!
commit
```

```
PEy(config)#
router ospf 1
no redistribute isis
redistribute isis level-2 subnets route-map ISISintoOSPF
```

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

Task 3: Modifying Administrative Distance

Step 1 Modify administrative distance on the PE routers:

```
RP/0/RSP0/CPU0:PEx(config)#
router ospf 1
distance ospf external 116
!
commit
```

```
PEy(config)#  
router ospf 1  
distance ospf external 116
```

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

Task 4: One-Way Redistribution

Step 1 Remove IS-IS into OSPF redistribution on the PE routers:

```
RP/0/RSP0/CPU0: PEx(config)#  
router ospf 1  
no redistribute isis 1  
!  
commit
```

```
PEy(config)#  
router ospf 1  
no redistribute isis
```

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

Step 2 Originate the OSPF default route on the PE routers:

```
RP/0/RSP0/CPU0: PEx(config)#  
router ospf 1  
default-information originate always  
!  
commit
```

```
PEy(config)#  
router ospf 1  
default-information originate always
```

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

Hardware Lab 6: Influence BGP Route Selection

Task 1: Setting the Second Link between the CE and PE Routers

Step 1 Enable the second link between the CE and PE routers:

```
SWx(config)#
vlan y0
!
interface FastEthernet0/2
switchport mode trunk
no shutdown
!
interface FastEthernet0/23
switchport access vlan y0
no shutdown
```

```
SWy(config)#
vlan x0
!
interface FastEthernet0/2
switchport mode trunk
no shutdown
!
interface FastEthernet0/23
switchport access vlan x0
no shutdown
```

```
SWxy(config)#
vlan x0,y0
!
interface FastEthernet0/1
port-type nni
switchport access vlan x0
no shutdown
!
interface FastEthernet0/2
port-type nni
switchport access vlan y0
no shutdown
!
interface FastEthernet0/21
port-type nni
switchport access vlan y0
no shutdown
!
interface FastEthernet0/23
port-type nni
switchport access vlan x0
no shutdown
```

```
CEx(config)#
interface GigabitEthernet0/1
ip address 192.168.20x.x1 255.255.255.0
no shutdown
speed 100
```

```
RP/0/RSP0/CPU0: PEx(config)#
interface GigabitEthernet0/0/0/0.y
  ipv4 address 192.168.20y.x0 255.255.255.0
  encapsulation dot1q y0
!
commit
```

```
CEy(config)#
interface GigabitEthernet0/1
  ip address 192.168.20y.y1 255.255.255.0
  no shutdown
  speed 100
```

```
PEy(config)#
interface GigabitEthernet0/0/0.x
  encapsulation dot1Q x0
  ip address 192.168.20x.y0 255.255.255.0
```

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

Step 2 Disable OSPF between the CE and PE routers:

```
CEx(config)#
no router ospf 1
```

```
RP/0/RSP0/CPU0: PEx(config)#
no router ospf 1
!
commit
```

```
CEy(config)#
no router ospf 1
```

```
PEy(config)#
no router ospf 1
```

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

Task 2: Configuring BGP Weight and Local Preference

Step 2 Enable a second EBGP session between the CE and PE routers.

```
CEx(config)#
router bgp 6450x
  neighbor 192.168.20x.y0 remote-as 64500
```

```
RP/0/RSP0/CPU0:PEx(config)#
router bgp 64500
 neighbor 192.168.20y.y1
  remote-as 6450y
  address-family ipv4 unicast
  route-policy pass out
  route-policy pass in
!
commit
```

```
CEy(config)#
router bgp 6450y
 neighbor 192.168.20y.x0 remote-as 64500
```

```
PEy(config)#
router bgp 64500
 neighbor 192.168.20x.x1 remote-as 6450x
```

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

Step 3 Configure BGP weight:

```
CEx(config)#
router bgp 6450x
 neighbor 192.168.10x.x0 weight 1
 neighbor 192.168.20x.y0 weight 2
```

```
CEy(config)#
router bgp 6450y
 neighbor 192.168.10y.y0 weight 1
 neighbor 192.168.20y.x0 weight 2
```

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

Step 4 Configure BGP local preference:

```
RP/0/RSP0/CPU0:PEx(config)#
route-policy LP99
  set local-preference 99
end-policy
!
router bgp 64500
 neighbor 192.168.10x.x1
  address-family ipv4 unicast
  route-policy LP99 in
!
commit
```

```

PEy(config)#
route-map LP99 permit
  set local-preference 99
!
router bgp 64500
  neighbor 192.168.10y.y1 route-map LP99 in

```

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

Task 3: Configuring BGP Multiexit Discriminator

Step 1 Remove BGP local preference configuration:

```

RP/0/RSP0/CPU0: PEx(config)#
router bgp 64500
  neighbor 192.168.10x.x1
    address-family ipv4 unicast
      no route-policy LP99 in
      route-policy pass in
  !
commit

```

```

PEy(config)#
router bgp 64500
  no neighbor 192.168.10y.y1 route-map LP99 in

```

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

Step 3 Send the BGP MED to the BGP neighbors:

```

CEx(config)#
route-map MED1000 permit
  set metric 1000
!
route-map MED1001 permit
  set metric 1001
!
router bgp 6450x
  neighbor 192.168.10x.x0 route-map MED1001 out
  neighbor 192.168.20x.y0 route-map MED1000 out

```

```

CEy(config)#
route-map MED1000 permit
  set metric 1000
!
route-map MED1001 permit
  set metric 1001
!
router bgp 6450y
  neighbor 192.168.10y.y0 route-map MED1001 out
  neighbor 192.168.20y.x0 route-map MED1000 out

```

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

Task 4: Configuring the BGP Community

Step 1 Remove the route-map:

```
CEx(config)#
router bgp 6450x
no neighbor 192.168.10x.x0 route-map MED1001 out
no neighbor 192.168.20x.y0 route-map MED1000 out
```

```
CEy(config)#
router bgp 6450y
no neighbor 192.168.10y.y0 route-map MED1001 out
no neighbor 192.168.20y.x0 route-map MED1000 out
```

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

Step 3 Configure the BGP community policy:

```
RP/0/RSP0/CPU0:PEx(config)#
route-policy from_Cust
  if community matches-any (64500:101) then
    set local-preference 101
  endif
  if community matches-any (64500:99) then
    set local-preference 99
  endif
  if community matches-any (64500:999) then
    set med 999
  endif
  if community matches-any (64500:1000) then
    set med 1000
  endif
  if community matches-any (64500:1001) then
    set med 1001
  else
    pass
  endif
end-policy
!
route-policy to_Cust
  if community matches-any (64500:1) then
    prepend as-path 64500
  endif
  if community matches-any (64500:2) then
    prepend as-path 64500 2
  endif
  if community matches-any (64500:3) then
    prepend as-path 64500 3
  else
    pass
  endif
end-policy
!
commit
```

```

PEy(config)#
ip community-list expanded c101 permit 64500:101_
ip community-list expanded c99 permit 64500:99_
ip community-list expanded c999 permit 64500:999_
ip community-list expanded c1000 permit 64500:1000_
ip community-list expanded c1001 permit 64500:1001_
ip community-list expanded c1 permit 64500:1_
ip community-list expanded c2 permit 64500:2_
ip community-list expanded c3 permit 64500:3_
!
route-map from_Cust permit 10
  match community c101
  set local-preference 101
route-map from_Cust permit 20
  match community c99
  set local-preference 99
route-map from_Cust permit 30
  match community c999
  set metric 999
route-map from_Cust permit 40
  match community c1000
  set metric 1000
route-map from_Cust permit 50
  match community c1001
  set metric 1001
route-map from_Cust permit 60
!
route-map to_Cust permit 10
  match community c1
  set as-path prepend 64500
route-map to_Cust permit 20
  match community c2
  set as-path prepend 64500 64500
route-map to_Cust permit 30
  match community c3
  set as-path prepend 64500 64500 64500
route-map to_Cust permit 40

```

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

Step 4 Apply the RPL and route-map:

```

RP/0/RSP0/CPU0:PEx(config)#
router bgp 64500
  neighbor 192.168.10x.x1
    address-family ipv4 unicast
      route-policy to_Cust out
      route-policy from_Cust in
  neighbor 192.168.20y.y1
    address-family ipv4 unicast
      route-policy to_Cust out
      route-policy from_Cust in
!
commit

```

```

PEy(config)#
router bgp 64500
neighbor 192.168.20x.x1 route-map to_Cust out
neighbor 192.168.10y.y1 route-map to_Cust out
neighbor 192.168.20x.x1 route-map from_Cust in
neighbor 192.168.10y.y1 route-map from_Cust in

```

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

Step 5 Enable BGP community support:

```

CEx(config)#
ip bgp-community new-format
!
router bgp 6450x
neighbor 192.168.10x.x0 send-community both
neighbor 192.168.20x.y0 send-community both

```

```

CEy(config)#
ip bgp-community new-format
!
router bgp 6450y
neighbor 192.168.10y.y0 send-community both
neighbor 192.168.20y.x0 send-community both

```

```

PEy(config)#
ip bgp-community new-format
!
router bgp 64500
neighbor 10.0.1.1 send-community both

```

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

Step 6 Set the BGP communities to change BGP local preference:

```

CEx(config)#
route-map Primary_SP permit 10
set community 64500:101
route-map Backup_SP permit 10
set community 64500:99
!
router bgp 6450x
neighbor 192.168.10x.x0 route-map Backup_SP out
neighbor 192.168.20x.y0 route-map Primary_SP out

```

```

CEy(config)#
route-map Primary_SP permit 10
  set community 64500:101
route-map Backup_SP permit 10
  set community 64500:99
!
router bgp 6450y
neighbor 192.168.10y.y0 route-map Backup_SP out
neighbor 192.168.20y.x0 route-map Primary_SP out

```

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

Step 8 Set the BGP communities to change the BGP MED:

```

CEx(config)#
no route-map Primary_SP permit 10
route-map Primary_SP permit 10
  set community 64500:999
no route-map Backup_SP permit 10
route-map Backup_SP permit 10
  set community 64500:1000

```

```

CEy(config)#
no route-map Primary_SP permit 10
route-map Primary_SP permit 10
  set community 64500:999
no route-map Backup_SP permit 10
route-map Backup_SP permit 10
  set community 64500:1000

```

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

Step 10 Set the BGP communities to change the BGP AS-path prepend:

```

CEx(config)#
no route-map Primary_SP permit 10
route-map Primary_SP permit 10
  set community 64500:1
no route-map Backup_SP permit 10
route-map Backup_SP permit 10
  set community 64500:3

```

```

CEy(config)#
no route-map Primary_SP permit 10
route-map Primary_SP permit 10
  set community 64500:1
no route-map Backup_SP permit 10
route-map Backup_SP permit 10
  set community 64500:3

```

x should be substituted with 1 or 3, and y should be substituted with 2 or 4, depending on your team that is assigned.

