

Deploying Cisco Service Provider Network Routing

Version 1.01

Lab Guide

Text Part Number: 97-3149-02



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.

Table of Contents

Lab Guide	1
Overview	1
Outline	1
Job Aids	2
Pod Access Information	2
Device Information	2
IP Addressing	4
NET Addressing	6
Lab 2-1: Implement OSPF Routing	7
Activity Objective	7
Visual Objective	7
OSPF Areas	8
Required Resources	8
Command List	9
Task 1: Enable OSPF on the Routers	11
Task 2: Influence OSPF DR and BDR Election	23
Task 3: Influence OSPF Route Selection by Changing OSPF Link Cost	23
Task 4: Configure OSPF Authentication	25
Task 5: Configure OSPF Virtual Links	26
Lab 2-2: Implement OSPF Special Area Types	31
Activity Objective	31
Visual Objective	31
OSPF Areas	32
Required Resources	32
Command List	33
Task 1: Enable OSPF Summarization	35
Task 2: Configure OSPF Stub Area	36
Lab 3-1: Implement Integrated IS-IS Routing	41
Activity Objective	41
Visual Objective	41
NET Addressing	42
Required Resources	42
Command List	43
Task 1: Enable Integrated IS-IS on the Routers	45
Task 2: IS-IS Route Summarization	48
Task 3: Enable IPv6 IS-IS Routing	49
Lab 4-1: Implement Basic BGP Routing	51
Activity Objective	51
Visual Objective	51
NET Addressing	52
Required Resources	52
Command List	53
Task 1: Configure External BGP	56
Task 2: Configure Internal BGP	57
Lab 5-1: Implement Route Redistribution	59
Activity Objective	59
Visual Objective	59
Required Resources	60
Command List	61
Task 1: Redistribute Between OSPF and IS-IS	63
Task 2: Prevent Potential Routing Loops	66
Task 3: Modify Administrative Distance	69
Task 4: One-way redistribution	72
Lab 5-2: Influence BGP Route Selection	74
Activity Objective	74
Visual Objective	74
Required Resources	75
Command List	76

Task 0: Set Second Link Between CE and PE	78
Task 1: Configure BGP Weight and Local Preference	79
Task 2: Configure BGP Multi-Exit-Discriminator	80
Task 3: Configure BGP Community	82
Answer Key	87
Lab 2-1 Answer Key: Implement OSPF Routing	87
Lab 2-2 Answer Key: Implement OSPF Special Area Types	93
Lab 3-1 Answer Key: Implement Integrated IS-IS Routing	96
Lab 4-1 Answer Key: Implement Basic BGP Routing	98
Lab 5-1 Answer Key: Implement Route Redistribution	100
Lab 5-2 Answer Key: Influence BGP Route Selection	103
Appendix A: Lab Topology	110

Lab Guide

Overview

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

Outline

This guide includes these activities:

- Job Aids
- Lab 2-1: Implement OSPF Routing
- Lab 2-2: Implement OSPF Special Area Types
- Lab 3-1: Implement Integrated IS-IS Routing
- Lab 4-1: Implement Basic BGP Routing
- Lab 5-1: Implement Route Redistribution
- Lab 5-2: Influence BGP Route Selection
- Appendix A: Tear-Out

Job Aids

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

Pod Access Information

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


Device Information

This lab topology consists of four (4) teams and eight (8) pods. Two students will work in one pod and two pods will work in one team. Each pod has one switch and two routers. Two pods share one additional switch. All teams share 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, 3, 5, or 7) will work on the PE router running Cisco IOS XR Software, and the second pod within the same team (pods 2, 4, 6, or 8) 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:

SPROUTE Lab Topology




Device Roles and Loopback IP Addresses

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

The following figure illustrates the interface identification used in this lab setup.


SPROUTE Lab Interface Identification: Team View



IP Addressing


The following figure illustrates the IP addressing scheme used in this lab setup.

SPROUTE Lab IP Addressing: Team View



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

SPROUTE Management IP Addresses



Note Replace the x or y with your pod number to get the IP addresses within your pod (x is for odd number pods 1, 3, 5, and 7; y is for even number pods 2, 4, 6, and 8). Replace the xy (where x < y) with numbers of the pods within the same team (for example, 12, 34, 56, or 78) to get 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
PE6	POS0/2/0	192.168.221.60/24	2001:db8:192:168:221::60/80
	POS0/2/1	192.168.222.60/24	2001:db8:192:168:222::60/80
PE8	POS0/2/0	192.168.221.80/24	2001:db8:192:168:221::80/80
	POS0/2/1	192.168.222.80/24	2001:db8:192:168:222::80/80
PEx	Gi0/0/0/0	192.168.10x.x0/24	2001:db8:192:168:10x::x0/80
	Gi0/0/0/1	192.168.1xy.x0/24	2001:db8:192:168:1xy::x0/80

Device	Interface	IPv4 Address	IPv6 Address
PEy	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
	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

NET Addressing

This subtopic includes table with NET addresses used in the lab activities.

Pod and Backbone NET Addresses

Replace the x or y with your pod number to get the NET addresses for routers.

Router	NET address
P1	49.0000.0100.0000.1001.00
P2	49.0000.0100.0000.2001.00
CEx	49.000x.0100.0x01.0001.00
CEy	49.000y.0100.0y01.0001.00
PEx	49.000x.0100.0x00.1001.00
PEy	49.000y.0100.0y00.1001.00

Lab 2-1: Implement OSPF Routing

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

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 in 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 used in the lab. Replace the x or y with your pod number to get the OSPF areas 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 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
<code>[no] shutdown</code>	Enables or disables interface
<code>area area virtual-link IP-address</code>	(Router) Enables OSPF virtual link
<code>clear ip ospf process</code>	Clears OSPF processes on the router
<code>configure terminal</code>	Enters configuration mode
<code>interface interface</code>	Enters interface configuration mode
<code>ip ospf authentication message-digest</code>	Enables OSPF MD5 authentication on the interface
<code>ip ospf cost cost</code>	Sets OSPF link cost on the interface
<code>ip ospf message-digest-key key-ID md5 key</code>	Sets OSPF MD5 key on the interface
<code>ip ospf priority priority</code>	Sets OSPF interface priority for DR and BDR election
<code>ip ipv6 address ip_address subnet_mask</code>	Sets an IPv4 or IPv6 address and the subnet mask on the interface
<code>ipv6 enable</code>	Enables IPv6 support on the interface
<code>ipv6 ospf process-id area area</code>	Enables OSPFv3 routing on the interface
<code>ipv6 unicast-routing</code>	Globally enables IPv6 unicast routing
<code>network prefix wildcard_mask area area</code>	(Router) Enables OSPF routing on the network and places network into OSPF area
<code>ping dest_IP source source_IP</code>	Verifies connectivity between source IP and destination IP
<code>router ospf process-id</code>	Enables OSPFv2 process on the router
<code>router-id IP-address</code>	(Router) Sets OSPF router ID
<code>show ip interface brief</code>	Displays interface status and IPv4 addresses configured
<code>show ip ospf</code>	Displays information related to the OSPF routing protocol running on the router
<code>show ip ospf interface</code>	Displays OSPF interface information
<code>show ip ospf virtual-links</code>	Displays OSPF virtual link information
<code>show ip ipv6 ospf database</code>	Displays the content of the OSPF database
<code>show ip ipv6 ospf neighbors</code>	Displays OSPF neighbor information
<code>show ip ipv6 protocols</code>	Displays IPv4 or IPv6 protocols running on the router
<code>show ip ipv6 route</code>	Displays the current routes in the routing table
<code>show running-config</code>	Displays running configuration

Cisco IOS XR Commands

Command	Description
<code>[no] shutdown</code>	Enables or disables interface
<code>area area</code>	(Router) Creates OSPF area on the router
<code>authentication message-digest</code>	(Router) Enables OSPF MD5 authentication on the interface
<code>clear ospf process</code>	Clears OSPF processes on the router
<code>commit</code>	Commits changes to the running configuration
<code>configure terminal</code>	Enters configuration mode
<code>cost cost</code>	(Router) Sets OSPF link cost on the interface
<code>interface interface</code>	(Global) Enters interface configuration mode
<code>interface interface</code>	(Router) Defines the interfaces on which the OSPF protocol runs
<code>ipv4 ipv6 address ip_address/len</code>	Sets the IPv4 or IPv6 address for an interface and the subnet mask using the prefix length format
<code>ipv6 enable</code>	Enables IPv6 support on the interface
<code>message-digest-key key-ID md5 key</code>	(Router) Sets OSPF MD5 key on the interface
<code>ping dest_IP source [source_IP]</code>	Verifies connectivity between source IP and destination IP (IPv4 and IPv6)
<code>priority priority</code>	(Router) Sets OSPF interface priority for DR and BDR election
<code>router ospf ospfv3 process-ID</code>	Creates a OSPFv2 or OSPFv3 process
<code>router-id IP-address</code>	(Router) Sets OSPF router ID.
<code>show ipv4 interface brief</code>	Displays interface status and IPv4 addresses configured
<code>show ospf</code>	Displays information related to the OSPF routing protocol running on the router
<code>show ospf interface</code>	Displays OSPF interface information
<code>show ospf ospfv3 database</code>	Displays the content of the OSPF database
<code>show ospf ospfv3 neighbors</code>	Displays OSPF neighbor information
<code>show protocols</code>	Displays protocols running on the router
<code>show route</code>	Displays the current routes in the routing table
<code>show running-config</code>	Displays running configuration
<code>traceroute IP-address</code>	Traces IP address
<code>virtual-link IP-address</code>	(Router) Enables OSPF virtual link

Task 1: Enable 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.

Activity Procedure

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

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

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

```
CE1#show ip interface brief | include GigabitEthernet0/0|Loopback0
GigabitEthernet0/0      192.168.101.11  YES manual up
Loopback0              10.1.10.1      YES manual up
```

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

```
RP/0/RSP0/CPU0:PE1#show ipv4 interface brief | include Up
Thu Jun  8 00:08:39.335 UTC
Loopback0              10.1.1.1        Up
MgmtEth0/RSP0/CPU0/0   10.10.10.17     Up
GigabitEthernet0/0/0/0  192.168.101.10  Up
GigabitEthernet0/0/0/1  192.168.112.10  Up
GigabitEthernet0/0/0/2  192.168.11.10   Up
GigabitEthernet0/0/0/3  192.168.12.10   Up
```

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

Successful ping from PE to CE router:

```
RP/0/RSP0/CPU0:PE1#ping 192.168.101.11
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.101.11, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

Successful ping between PE routers in your team:

```
RP/0/RSP0/CPU0:PE1#ping 192.168.112.20
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.112.20, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

Successful ping between PE router in your team and core routers P1 and P2:

```
RP/0/RSP0/CPU0:PE1#ping 192.168.11.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.11.1, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
RP/0/RSP0/CPU0:PE1#ping 192.168.12.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.12.2, 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 OSPF routing in your team routers:

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

Link (Interface)	OSPF Area
PEx (Gi0/0/0/2) – P1	
PEx (Gi0/0/0/3) – P2	
PEx (Gi0/0/0/1) - PEy (Gi0/0/0/1)	
PEx (Gi0/0/0/0) – CEx (Gi0/0)	
PEx (Loopback0)	
CEx (Loopback0)	

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

- Step 5** You should be able to successfully test IP connectivity between CE routers in your team.

Successful ping from CE router in your pod and CE router in neighbor pod (same team).

```
CE1#ping 10.2.10.1
```

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

```
Sending 5, 100-byte ICMP Echos to 10.2.10.1, timeout is 2 seconds:
```

```
!!!!
```

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

- Step 6** Verify what router ID is assigned to the OSPF routing process on CE and PE routers. Find out what OSPF router IDs are assigned in the neighbor pod (same team).

The OSPF router ID on the CE and PE routers:

```
CE1#show ip ospf | include ID
Routing Process "ospf 1" with ID 10.1.10.1
```

```
RP/0/RSP0/CPU0:PE1#show ospf | include ID
Routing Process "ospf 1" with ID 10.1.1.1
```

Write the results into the following table:

Router	OSPF Router ID
PEx (pod)	
CEx (pod)	
PEy (neighbor pod)	
CEy (neighbor pod)	

- Step 7** On the CE and PE routers in your pod, set OSPF router IDs to the IP addresses x0.x0.x0.x0 (for CEx) and x.x.x.x (for PEx) or y0.y0.y0.y0 (for CEy) and y.y.y.y (for PEy), where x or y is your pod number. 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 Gi0/0 and Lo0 interfaces and configure IPv6 address on the Lo0 interface only. On the PE router, enable IPv6 on Gi0/0/0/0, Gi0/0/0/1, Gi0/0/0/2, Gi0/0/0/3, and Lo0 interfaces and configure IP address on the Lo0 interface only. Find correct IPv6 addresses in the visual objective figure.

- 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.

Activity Verification

You have completed this task when you attain these results:

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

```
CE1#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 10.10.10.10
  Number of areas in this router is 1. 1 normal 0 stub 0 nssa
  Maximum path: 4
  Routing for Networks:
    10.1.10.1 0.0.0.0 area 11
    192.168.101.11 0.0.0.0 area 11
  Routing Information Sources:
    Gateway         Distance      Last Update
    1.1.1.1          110           00:07:12
    10.1.1.1         110           00:13:14
  Distance: (default is 110)
```

```
CE1#show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
1.1.1.1	1	FULL/BDR	00:00:38	192.168.101.10	Gi0/0

```
CE1#show ip ospf database
```

OSPF Router with ID (10.10.10.10) (Process ID 1)

Router Link States (Area 11)

Link ID	ADV Router	Age	Seq#	Checksum	Link count
1.1.1.1	1.1.1.1	558	0x80000003	0x000B6A	1
10.10.10.10	10.10.10.10	550	0x80000003	0x0008FD	2

Net Link States (Area 11)

Link ID	ADV Router	Age	Seq#	Checksum
192.168.101.11	10.10.10.10	550	0x80000001	0x007E8D

Summary Net Link States (Area 11)

Link ID	ADV Router	Age	Seq#	Checksum
10.0.1.1	1.1.1.1	558	0x80000002	0x00E544
10.0.2.1	1.1.1.1	558	0x80000002	0x00DA4E
10.1.1.1	1.1.1.1	559	0x80000002	0x00CF5A
10.2.1.1	1.1.1.1	482	0x80000005	0x00C75D
10.2.10.1	1.1.1.1	447	0x80000001	0x0076A8
192.168.1.0	1.1.1.1	558	0x80000002	0x00C00B
192.168.2.0	1.1.1.1	558	0x80000002	0x00B515
192.168.11.0	1.1.1.1	559	0x80000002	0x00487A
192.168.12.0	1.1.1.1	558	0x80000005	0x003787
192.168.21.0	1.1.1.1	559	0x80000002	0x00E3D3
192.168.22.0	1.1.1.1	559	0x80000002	0x00D8DD
192.168.31.0	1.1.1.1	559	0x80000002	0x007538
192.168.32.0	1.1.1.1	559	0x80000002	0x006A42
192.168.41.0	1.1.1.1	559	0x80000002	0x00079C
192.168.42.0	1.1.1.1	560	0x80000002	0x00FBA6
192.168.51.0	1.1.1.1	560	0x80000002	0x009801
192.168.52.0	1.1.1.1	560	0x80000002	0x008D0B

192.168.61.0	1.1.1.1	560	0x80000002	0x002A65
192.168.62.0	1.1.1.1	560	0x80000002	0x001F6F
192.168.71.0	1.1.1.1	560	0x80000002	0x00BBC9
192.168.72.0	1.1.1.1	560	0x80000002	0x00B0D3
192.168.81.0	1.1.1.1	560	0x80000002	0x004D2E
192.168.82.0	1.1.1.1	560	0x80000002	0x004238
192.168.102.0	1.1.1.1	448	0x80000001	0x0067FF
192.168.112.0	1.1.1.1	493	0x80000006	0x00E474

- On the PE router in your pod, verify IP protocols running, OSPF neighbors, and the OSPF database. There should be four OSPF neighbors. The OSPF database should show LSAs from areas 0, x, and 1x, where “x” is your pod number. Outputs should be similar to the following:

```
RP/0/RSP0/CPU0:PE1#show protocols
```

```
Routing Protocol OSPF 1
  Router Id: 1.1.1.1
  Distance: 110
  Non-Stop Forwarding: Disabled
  Redistribution:
    None
  Area 0
    GigabitEthernet0/0/0/1
    GigabitEthernet0/0/0/2
    GigabitEthernet0/0/0/3
  Area 1
    Loopback0
  Area 11
    GigabitEthernet0/0/0/0
```

```
% EIGRP: EIGRP not configured
```

```
RP/0/RSP0/CPU0:PE1#show ospf neighbor
```

```
* Indicates MADJ interface
```

```
Neighbors for OSPF 1
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
2.2.2.2	1	FULL/BDR	00:00:37	192.168.112.20	Gi0/0/0/1
Neighbor is up for 00:56:16					
0.0.0.1	1	FULL/DR	00:00:34	192.168.11.1	Gi0/0/0/2
Neighbor is up for 00:57:31					
0.0.0.2	1	FULL/DR	00:00:33	192.168.12.2	Gi0/0/0/3
Neighbor is up for 00:57:32					
10.10.10.10	1	FULL/DR	00:00:38	192.168.101.11	Gi0/0/0/0
Neighbor is up for 00:57:33					

```
Total neighbor count: 4
```

```
RP/0/RSP0/CPU0:PE1#show ospf database
```

```
OSPF Router with ID (1.1.1.1) (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	559	0x80000c12	0x00756a	7
0.0.0.2	0.0.0.2	561	0x80000c0c	0x00a364	8
1.1.1.1	1.1.1.1	562	0x80000008	0x00321e	3
2.2.2.2	2.2.2.2	557	0x80000003	0x00679f	3
3.3.3.3	3.3.3.3	617	0x8000000a	0x006493	4
4.4.4.4	4.4.4.4	613	0x80000002	0x00ce78	2
5.5.5.5	5.5.5.5	856	0x80000004	0x000bf8	3
10.7.1.1	10.7.1.1	1137	0x800003cc	0x00bb21	3
10.7.10.1	10.7.10.1	1795	0x800003c0	0x00ca09	2
10.8.10.1	10.8.10.1	696	0x800005ce	0x00c2e6	2
10.100.100.100	10.100.100.100	670	0x800003d3	0x003516	3
30.30.30.30	30.30.30.30	4	(DNA) 0x80000002	0x00a49d	1

Net Link States (Area 0)

Link ID	ADV Router	Age	Seq#	Checksum
192.168.2.2	0.0.0.2	581	0x80000b7e	0x00c17d
192.168.11.1	0.0.0.1	592	0x80000005	0x009822
192.168.12.2	0.0.0.2	599	0x80000005	0x00872f
192.168.21.1	0.0.0.1	559	0x80000003	0x00604e
192.168.22.2	0.0.0.2	561	0x80000003	0x004f5b
192.168.31.1	0.0.0.1	1863	0x80000006	0x001e7f
192.168.32.2	0.0.0.2	1847	0x80000006	0x000d8c
192.168.51.1	0.0.0.1	617	0x80000004	0x00a9d9
192.168.52.2	0.0.0.2	842	0x80000004	0x0098e6
192.168.71.70	10.7.1.1	1137	0x800003bd	0x00e477
192.168.82.80	10.100.100.100	670	0x800003bb	0x0025da
192.168.107.70	10.7.1.1	619	0x80000395	0x00a4a0
192.168.108.81	10.8.10.1	696	0x800003b8	0x005d89
192.168.112.10	1.1.1.1	562	0x80000001	0x001d29
192.168.134.30	3.3.3.3	617	0x80000003	0x00c938

Summary Net Link States (Area 0)

Link ID	ADV Router	Age	Seq#	Checksum
10.1.1.1	1.1.1.1	595	0x80000002	0x00cf5a
10.1.10.1	1.1.1.1	573	0x80000001	0x0078a8
10.2.1.1	2.2.2.2	558	0x80000001	0x00a77e
10.2.10.1	2.2.2.2	519	0x80000001	0x004ecd
10.3.1.1	3.3.3.3	1723	0x80000003	0x0079a5
10.3.10.1	3.3.3.3	892	0x80000003	0x0020f4
10.3.10.1	30.30.30.30	85 (DNA)	0x80000002	0x00eabf
10.4.1.1	4.4.4.4	608	0x80000001	0x0053c8
10.4.10.1	4.4.4.4	491	0x80000001	0x00f918
10.5.10.1	5.5.5.5	608	0x80000002	0x00cd3e
13.0.0.1	30.30.30.30	95 (DNA)	0x80000001	0x00585d
192.168.101.0	1.1.1.1	573	0x80000005	0x006005
192.168.102.0	2.2.2.2	558	0x80000001	0x003f25
192.168.103.0	3.3.3.3	1562	0x80000005	0x000e4d
192.168.103.0	30.30.30.30	85 (DNA)	0x80000002	0x00e60b
192.168.104.0	4.4.4.4	491	0x80000001	0x00ec6d
192.168.105.0	5.5.5.5	608	0x80000006	0x00b996

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	353	0x800000e1	0x00e55e	0
1.0.0.0	0.0.0.2	580	0x800000e1	0x00e859	0
1.0.0.0	10.7.1.1	619	0x800000e1	0x0084a6	0
1.0.0.0	10.100.100.100	671	0x800000e1	0x009570	0
1.0.0.4	10.100.100.100	670	0x800000e6	0x002a15	4
1.0.0.8	0.0.0.1	353	0x800000e6	0x00cdf9	8
1.0.0.9	10.7.1.1	1644	0x800000e6	0x005016	9
1.0.0.11	0.0.0.2	580	0x800000e8	0x006619	11
1.0.0.12	0.0.0.1	353	0x800000e8	0x004c2b	12
1.0.0.19	0.0.0.2	580	0x800000e6	0x00b94d	19

Router Link States (Area 1)

Link ID	ADV Router	Age	Seq#	Checksum	Link count
1.1.1.1	1.1.1.1	599	0x80000001	0x001309	1

Summary Net Link States (Area 1)

Link ID	ADV Router	Age	Seq#	Checksum
10.0.1.1	1.1.1.1	589	0x80000002	0x00e544
10.0.2.1	1.1.1.1	594	0x80000001	0x00dc4d
10.1.10.1	1.1.1.1	573	0x80000001	0x0078a8
10.2.1.1	1.1.1.1	557	0x80000001	0x00cf59
10.2.10.1	1.1.1.1	518	0x80000001	0x0076a8
10.3.1.1	1.1.1.1	589	0x80000002	0x00cb5a

10.3.10.1	1.1.1.1	589	0x80000002	0x0072a9
10.4.1.1	1.1.1.1	589	0x80000002	0x00c95a
10.4.10.1	1.1.1.1	488	0x80000001	0x0072a8
10.5.10.1	1.1.1.1	589	0x80000002	0x005abf
10.7.1.1	1.1.1.1	589	0x80000002	0x00f523
10.7.10.1	1.1.1.1	589	0x80000002	0x009c72
10.8.1.1	1.1.1.1	594	0x80000001	0x00eb2d
10.8.10.1	1.1.1.1	594	0x80000001	0x00927c
13.0.0.1	1.1.1.1	589	0x80000002	0x00dd48
192.168.2.0	1.1.1.1	594	0x80000001	0x00c109
192.168.11.0	1.1.1.1	599	0x80000001	0x004a79
192.168.12.0	1.1.1.1	599	0x80000001	0x003f83
192.168.21.0	1.1.1.1	561	0x80000004	0x00dfd5
192.168.22.0	1.1.1.1	594	0x80000001	0x00dad5
192.168.31.0	1.1.1.1	589	0x80000002	0x007538
192.168.32.0	1.1.1.1	589	0x80000002	0x007437
192.168.41.0	1.1.1.1	589	0x80000002	0x001b86
192.168.42.0	1.1.1.1	594	0x80000001	0x00fda5
192.168.51.0	1.1.1.1	589	0x80000002	0x009801
192.168.52.0	1.1.1.1	589	0x80000002	0x0097ff
192.168.71.0	1.1.1.1	589	0x80000002	0x001666
192.168.82.0	1.1.1.1	594	0x80000001	0x009ed3
192.168.101.0	1.1.1.1	573	0x80000001	0x006801
192.168.102.0	1.1.1.1	557	0x80000001	0x0067ff
192.168.103.0	1.1.1.1	589	0x80000002	0x0064ff
192.168.104.0	1.1.1.1	488	0x80000001	0x0065fd
192.168.105.0	1.1.1.1	589	0x80000002	0x004e14
192.168.107.0	1.1.1.1	589	0x80000002	0x0092c4
192.168.108.0	1.1.1.1	594	0x80000001	0x0089cd
192.168.112.0	1.1.1.1	567	0x80000003	0x00ea71
192.168.134.0	1.1.1.1	589	0x80000002	0x000e37
192.168.156.0	1.1.1.1	589	0x80000002	0x001b14

Router Link States (Area 11)

Link ID	ADV Router	Age	Seq#	Checksum	Link count
1.1.1.1	1.1.1.1	595	0x80000003	0x000b6a	1
10.10.10.10	10.10.10.10	574	0x80000003	0x0008fd	2

Net Link States (Area 11)

Link ID	ADV Router	Age	Seq#	Checksum
192.168.101.11	10.10.10.10	574	0x80000001	0x007e8d

Summary Net Link States (Area 11)

Link ID	ADV Router	Age	Seq#	Checksum
10.0.1.1	1.1.1.1	589	0x80000002	0x00e544
10.0.2.1	1.1.1.1	594	0x80000003	0x00d84f
10.1.1.1	1.1.1.1	594	0x80000002	0x00cf5a
10.2.1.1	1.1.1.1	557	0x80000006	0x00c55e
10.2.10.1	1.1.1.1	518	0x80000001	0x0076a8
10.3.1.1	1.1.1.1	589	0x80000002	0x00cb5a
10.3.10.1	1.1.1.1	589	0x80000002	0x0072a9
10.4.1.1	1.1.1.1	589	0x80000002	0x00c95a
10.4.10.1	1.1.1.1	488	0x80000001	0x0072a8
10.5.10.1	1.1.1.1	589	0x80000002	0x005abf
10.7.1.1	1.1.1.1	589	0x80000003	0x00f324
10.7.10.1	1.1.1.1	589	0x80000002	0x009c72
10.8.1.1	1.1.1.1	594	0x80000003	0x00e72f
10.8.10.1	1.1.1.1	594	0x80000003	0x008e7e
13.0.0.1	1.1.1.1	589	0x80000002	0x00dd48
192.168.2.0	1.1.1.1	594	0x80000002	0x00bf0a
192.168.11.0	1.1.1.1	594	0x80000004	0x00447c
192.168.12.0	1.1.1.1	594	0x80000004	0x003986
192.168.21.0	1.1.1.1	561	0x80000006	0x00dbd7
192.168.22.0	1.1.1.1	594	0x80000003	0x00d6de
192.168.31.0	1.1.1.1	589	0x80000002	0x007538

192.168.32.0	1.1.1.1	589	0x80000002	0x007437
192.168.41.0	1.1.1.1	589	0x80000002	0x001b86
192.168.42.0	1.1.1.1	594	0x80000003	0x00f9a7
192.168.51.0	1.1.1.1	589	0x80000002	0x009801
192.168.52.0	1.1.1.1	589	0x80000002	0x0097ff
192.168.71.0	1.1.1.1	589	0x80000002	0x001666
192.168.82.0	1.1.1.1	594	0x80000003	0x009ad5
192.168.102.0	1.1.1.1	557	0x80000001	0x0067ff
192.168.103.0	1.1.1.1	589	0x80000002	0x0064ff
192.168.104.0	1.1.1.1	488	0x80000001	0x0065fd
192.168.105.0	1.1.1.1	589	0x80000002	0x004e14
192.168.107.0	1.1.1.1	589	0x80000002	0x0092c4
192.168.108.0	1.1.1.1	594	0x80000003	0x0085cf
192.168.112.0	1.1.1.1	567	0x80000007	0x00e275
192.168.134.0	1.1.1.1	589	0x80000002	0x000e37
192.168.156.0	1.1.1.1	589	0x80000002	0x001b14

- 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:

```
CE1#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, + - replicated route
```

Gateway of last resort is not set

```

      10.0.0.0/32 is subnetted, 6 subnets
O IA      10.0.1.1 [110/3] via 192.168.101.10, 00:04:52, GigabitEthernet0/0
O IA      10.0.2.1 [110/3] via 192.168.101.10, 00:04:52, GigabitEthernet0/0
O IA      10.1.1.1 [110/2] via 192.168.101.10, 00:04:52, GigabitEthernet0/0
O IA      10.2.1.1 [110/3] via 192.168.101.10, 00:04:52, GigabitEthernet0/0
O IA      10.2.10.1 [110/4] via 192.168.101.10, 00:04:52, GigabitEthernet0/0
O IA 192.168.1.0/24 [110/3] via 192.168.101.10, 00:04:52, GigabitEthernet0/0
O IA 192.168.2.0/24 [110/3] via 192.168.101.10, 00:04:53, GigabitEthernet0/0
O IA 192.168.11.0/24 [110/2] via 192.168.101.10, 00:04:53, GigabitEthernet0/0
O IA 192.168.12.0/24 [110/2] via 192.168.101.10, 00:04:53, GigabitEthernet0/0
O IA 192.168.21.0/24 [110/3] via 192.168.101.10, 00:04:53, GigabitEthernet0/0
O IA 192.168.22.0/24 [110/3] via 192.168.101.10, 00:04:53, GigabitEthernet0/0
O IA 192.168.31.0/24 [110/3] via 192.168.101.10, 00:04:53, GigabitEthernet0/0
O IA 192.168.32.0/24 [110/3] via 192.168.101.10, 00:04:55, GigabitEthernet0/0
O IA 192.168.41.0/24 [110/3] via 192.168.101.10, 00:04:55, GigabitEthernet0/0
O IA 192.168.42.0/24 [110/3] via 192.168.101.10, 00:04:55, GigabitEthernet0/0
O IA 192.168.51.0/24 [110/3] via 192.168.101.10, 00:04:55, GigabitEthernet0/0
O IA 192.168.52.0/24 [110/3] via 192.168.101.10, 00:04:55, GigabitEthernet0/0
O IA 192.168.61.0/24 [110/3] via 192.168.101.10, 00:04:55, GigabitEthernet0/0
O IA 192.168.62.0/24 [110/3] via 192.168.101.10, 00:04:55, GigabitEthernet0/0
O IA 192.168.71.0/24 [110/3] via 192.168.101.10, 00:04:55, GigabitEthernet0/0
O IA 192.168.72.0/24 [110/3] via 192.168.101.10, 00:04:55, GigabitEthernet0/0
O IA 192.168.81.0/24 [110/3] via 192.168.101.10, 00:04:55, GigabitEthernet0/0
O IA 192.168.82.0/24 [110/3] via 192.168.101.10, 00:04:55, GigabitEthernet0/0
O IA 192.168.102.0/24
      [110/3] via 192.168.101.10, 00:04:55, GigabitEthernet0/0
O IA 192.168.112.0/24
      [110/2] via 192.168.101.10, 00:04:56, GigabitEthernet0/0
```

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

```
RP/0/RSP0/CPU0:PE1#show route ospf
```

```
Thu Jun  8 02:19:25.000 UTC
```

```

O      10.0.1.1/32 [110/2] via 192.168.11.1, 01:01:08, GigabitEthernet0/0/0/2
O      10.0.2.1/32 [110/2] via 192.168.12.2, 01:01:08, GigabitEthernet0/0/0/3
O      10.1.10.1/32 [110/2] via 192.168.101.11, 01:01:01, GigabitEthernet0/0/0/0
```

```

O IA 10.2.1.1/32 [110/2] via 192.168.112.20, 00:59:51, GigabitEthernet0/0/0/1
O IA 10.2.10.1/32 [110/3] via 192.168.112.20, 00:59:16, GigabitEthernet0/0/0/1
O 192.168.1.0/24 [110/2] via 192.168.11.1, 01:01:08, GigabitEthernet0/0/0/2
O 192.168.2.0/24 [110/2] via 192.168.11.1, 01:01:08, GigabitEthernet0/0/0/2
O 192.168.21.0/24 [110/2] via 192.168.11.1, 00:59:53,
GigabitEthernet0/0/0/2
[110/2] via 192.168.112.20, 00:59:53,
GigabitEthernet0/0/0/1
O 192.168.22.0/24 [110/2] via 192.168.12.2, 00:59:48,
GigabitEthernet0/0/0/3
[110/2] via 192.168.112.20, 00:59:48,
GigabitEthernet0/0/0/1
O 192.168.31.0/24 [110/2] via 192.168.11.1, 01:01:08,
GigabitEthernet0/0/0/2
O 192.168.32.0/24 [110/2] via 192.168.12.2, 01:01:08,
GigabitEthernet0/0/0/3
O 192.168.41.0/24 [110/2] via 192.168.11.1, 01:01:08,
GigabitEthernet0/0/0/2
O 192.168.42.0/24 [110/2] via 192.168.12.2, 01:01:08,
GigabitEthernet0/0/0/3
O 192.168.51.0/24 [110/2] via 192.168.11.1, 01:01:08,
GigabitEthernet0/0/0/2
O 192.168.52.0/24 [110/2] via 192.168.12.2, 01:01:08,
GigabitEthernet0/0/0/3
O 192.168.61.0/24 [110/2] via 192.168.11.1, 01:01:08,
GigabitEthernet0/0/0/2
O 192.168.62.0/24 [110/2] via 192.168.12.2, 01:01:08,
GigabitEthernet0/0/0/3
O 192.168.71.0/24 [110/2] via 192.168.11.1, 01:01:08,
GigabitEthernet0/0/0/2
O 192.168.72.0/24 [110/2] via 192.168.12.2, 01:01:08,
GigabitEthernet0/0/0/3
O 192.168.81.0/24 [110/2] via 192.168.11.1, 01:01:08,
GigabitEthernet0/0/0/2
O 192.168.82.0/24 [110/2] via 192.168.12.2, 01:01:08,
GigabitEthernet0/0/0/3
O IA 192.168.102.0/24 [110/2] via 192.168.112.20, 00:59:16,
GigabitEthernet0/0/0/1

```

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

CE1#**ping 10.2.10.1** (IP address of the CE2 router Lo0 interface)

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 10.2.10.1, timeout is 2 seconds:

!!!!

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

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

CE1#**show ipv6 protocols**

IPv6 Routing Protocol is "connected"

IPv6 Routing Protocol is "ND"

IPv6 Routing Protocol is "ospf 1"

Interfaces (Area 11):

Loopback0

GigabitEthernet0/0

Redistribution:

None

CE1#**show ipv6 ospf neighbor**

Neighbor ID	Pri	State	Dead Time	Interface ID	Interface
10.1.1.1	1	FULL/BDR	00:00:31	7	Gi0/0

```
CE1#show ipv6 ospf database
```

```
OSPFv3 Router with ID (11.0.0.1) (Process ID 1)
```

```
Router Link States (Area 11)
```

ADV Router	Age	Seq#	Fragment ID	Link count	Bits
10.1.1.1	621	0x80000002	0	1	B
11.0.0.1	619	0x80000002	0	1	None

```
Net Link States (Area 11)
```

ADV Router	Age	Seq#	Link ID	Rtr count
11.0.0.1	619	0x80000001	2	2

```
Inter Area Prefix Link States (Area 11)
```

ADV Router	Age	Seq#	Prefix
10.1.1.1	1228	0x80000001	2001:DB8:10:0:1::1/128
10.1.1.1	1228	0x80000001	2001:DB8:10:0:2::1/128
10.1.1.1	1228	0x80000001	2001:DB8:10:1:1::1/128
10.1.1.1	987	0x80000002	2001:DB8:10:2:1::1/128
10.1.1.1	723	0x80000001	2001:DB8:10:2:10::1/128

```
Link (Type-8) Link States (Area 11)
```

ADV Router	Age	Seq#	Link ID	Interface
10.1.1.1	1239	0x80000001	7	Gi0/0
11.0.0.1	1293	0x80000001	2	Gi0/0

```
Intra Area Prefix Link States (Area 11)
```

ADV Router	Age	Seq#	Link ID	Ref-lstype	Ref-LSID
11.0.0.1	1315	0x80000001	0	0x2001	0

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

```
RP/0/RSP0/CPU0:PE1#show protocols ipv6
```

```
% EIGRP: EIGRP not configured
```

```
Routing Protocol OSPFv3 1
```

```
Router Id: 10.1.1.1
```

```
Distance: 110
```

```
Graceful Restart: Disabled
```

```
Redistribution:
```

```
None
```

```
Area 0
```

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

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

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

```
Area 1
```

```
Loopback0
```

```
Area 11
```

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

```
RP/0/RSP0/CPU0:PE1#show ospfv3 neighbor
```

```
Neighbors for OSPFv3 1
```

Neighbor ID	Pri	State	Dead Time	Interface ID	Interface
10.1.10.1	1	FULL/DR	00:00:36	7	
GigabitEthernet0/0/0/0					
Neighbor is up for 00:10:21					
10.2.1.1	1	FULL/DR	00:00:37	8	
GigabitEthernet0/0/0/1					
Neighbor is up for 00:09:55					
10.0.1.1	1	FULL/DR	00:00:33	9	
GigabitEthernet0/0/0/2					

```

Neighbor is up for 00:10:21
10.0.2.1      1      FULL/DR      00:00:33      10
GigabitEthernet0/0/0/3
Neighbor is up for 00:10:21

```

Total neighbor count: 4

RP/0/RSP0/CPU0:PE1#show ospfv3 database

OSPFv3 Router with ID (10.1.1.1) (Process ID 1)

Router Link States (Area 0)

ADV Router	Age	Seq#	Fragment ID	Link count	Bits
10.0.1.1	625	0x80000b87	0	5	None
10.0.2.1	103	0x80000b9c	0	7	None
10.1.1.1	589	0x80000003	0	3	B
10.2.1.1	590	0x80000003	0	3	B
10.3.1.1	117	0x80000004	0	3	B
10.4.1.1	103	0x80000003	0	2	B
10.5.1.1	1815	0x80000002	0	2	B
10.100.100.100	933	0x800003c8	0	1	None

Net Link States (Area 0)

ADV Router	Age	Seq#	Link ID	Rtr count
10.0.1.1	642	0x80000001	9	2
10.0.1.1	1114	0x80000002	10	2
10.0.1.1	1820	0x80000001	11	2
10.0.1.1	625	0x80000001	13	2
10.0.2.1	927	0x80000b7b	11	2

RP/0/RSP0/CPU0:PE1(config-ospfv3-ar-if)#do show ospfv3 database
Tue Nov 15 19:20:01.918 UTC

OSPFv3 Router with ID (10.1.1.1) (Process ID 1)

Router Link States (Area 0)

ADV Router	Age	Seq#	Fragment ID	Link count	Bits
10.0.1.1	1048	0x80000b87	0	5	None
10.0.2.1	526	0x80000b9c	0	7	None
10.1.1.1	1012	0x80000003	0	3	B
10.2.1.1	1013	0x80000003	0	3	B
10.3.1.1	540	0x80000004	0	3	B
10.4.1.1	526	0x80000003	0	2	B
10.5.1.1	229	0x80000003	0	2	B
10.100.100.100	1356	0x800003c8	0	1	None

Net Link States (Area 0)

ADV Router	Age	Seq#	Link ID	Rtr count
10.0.1.1	1065	0x80000001	9	2
10.0.1.1	1537	0x80000002	10	2
10.0.1.1	293	0x80000002	11	2
10.0.1.1	1048	0x80000001	13	2
10.0.2.1	1350	0x80000b7b	11	2
10.0.2.1	1069	0x80000001	12	2
10.0.2.1	1600	0x80000002	13	2
10.0.2.1	361	0x80000002	14	2
10.0.2.1	1048	0x80000001	16	2
10.0.2.1	526	0x80000001	17	2
10.2.1.1	1013	0x80000001	9	2
10.3.1.1	540	0x80000001	8	2
10.100.100.100	1356	0x800003be	11	2

Inter Area Prefix Link States (Area 0)

ADV Router	Age	Seq#	Prefix
10.1.1.1	1059	0x80000001	2001:db8:10:1:1::1/128
10.1.1.1	1023	0x80000001	2001:db8:10:1:10::1/128

10.2.1.1	1044	0x80000001	2001:db8:10:2:1::1/128
10.2.1.1	1008	0x80000001	2001:db8:10:2:10::1/128
10.3.1.1	1545	0x80000002	2001:db8:10:3:1::1/128
10.3.1.1	1545	0x80000002	2001:db8:10:3:10::1/128
10.4.1.1	536	0x80000002	2001:db8:10:4:1::1/128
10.4.1.1	536	0x80000002	2001:db8:192:168:104::/80
10.4.1.1	466	0x80000001	2001:db8:10:4:10::1/128
10.4.1.1	456	0x80000001	2001:db8:192:168::/64
10.5.1.1	229	0x80000002	2001:db8:10:5:1::1/128
10.5.1.1	229	0x80000002	2001:db8:10:5:10::1/128

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

ADV Router	Age	Seq#	Link ID	Interface
10.1.1.1	1069	0x80000001	8	Gi0/0/0/1
10.2.1.1	1054	0x80000001	9	Gi0/0/0/1
10.0.1.1	534	0x80000002	9	Gi0/0/0/2
10.1.1.1	1069	0x80000001	9	Gi0/0/0/2
10.0.2.1	602	0x80000002	12	Gi0/0/0/3
10.1.1.1	1069	0x80000001	10	Gi0/0/0/3

Intra Area Prefix Link States (Area 0)

ADV Router	Age	Seq#	Link ID	Ref-lstype	Ref-LSID
10.0.1.1	793	0x80000b84	0	0x2001	0
10.0.2.1	1600	0x80000ba3	0	0x2001	0
10.0.2.1	526	0x80000001	17408	0x2002	17
10.3.1.1	540	0x80000001	8192	0x2002	8
10.100.100.100	1356	0x800003be	11264	0x2002	11

Router Link States (Area 1)

ADV Router	Age	Seq#	Fragment ID	Link count	Bits
10.1.1.1	1069	0x80000001	0	0	B

Inter Area Prefix Link States (Area 1)

ADV Router	Age	Seq#	Prefix
10.1.1.1	1059	0x80000001	2001:db8:192:168:82::/80
10.1.1.1	1059	0x80000001	2001:db8:10:0:1::1/128
10.1.1.1	1059	0x80000001	2001:db8:10:0:2::1/128
10.1.1.1	1059	0x80000001	2001:db8:10:5:10::1/128
10.1.1.1	1059	0x80000001	2001:db8:10:5:1::1/128
10.1.1.1	1059	0x80000001	2001:db8:10:3:10::1/128
10.1.1.1	1059	0x80000001	2001:db8:10:3:1::1/128
10.1.1.1	1008	0x80000002	2001:db8:10:2:1::1/128
10.1.1.1	1023	0x80000001	2001:db8:10:1:10::1/128
10.1.1.1	1007	0x80000001	2001:db8:10:2:10::1/128
10.1.1.1	533	0x80000001	2001:db8:192:168:134::/80
10.1.1.1	523	0x80000001	2001:db8:192:168:41::/80
10.1.1.1	523	0x80000001	2001:db8:10:4:1::1/128
10.1.1.1	523	0x80000001	2001:db8:192:168:104::/80
10.1.1.1	464	0x80000001	2001:db8:10:4:10::1/128
10.1.1.1	454	0x80000001	2001:db8:192:168::/64

Intra Area Prefix Link States (Area 1)

ADV Router	Age	Seq#	Link ID	Ref-lstype	Ref-LSID
10.1.1.1	1069	0x80000001	0	0x2001	0

Router Link States (Area 11)

ADV Router	Age	Seq#	Fragment ID	Link count	Bits
10.1.1.1	1030	0x80000002	0	1	B
10.1.10.1	1031	0x80000002	0	1	None

Net Link States (Area 11)

ADV Router	Age	Seq#	Link ID	Rtr count
10.1.10.1	1031	0x80000001	2	2

Inter Area Prefix Link States (Area 11)

ADV Router	Age	Seq#	Prefix
10.1.1.1	1059	0x80000001	2001:db8:192:168:82::/80
10.1.1.1	1059	0x80000001	2001:db8:10:0:1::1/128
10.1.1.1	1059	0x80000001	2001:db8:10:0:2::1/128
10.1.1.1	1059	0x80000001	2001:db8:10:1:1::1/128
10.1.1.1	1059	0x80000001	2001:db8:10:5:10::1/128
10.1.1.1	1059	0x80000001	2001:db8:10:5:1::1/128
10.1.1.1	1059	0x80000001	2001:db8:10:3:10::1/128
10.1.1.1	1059	0x80000001	2001:db8:10:3:1::1/128
10.1.1.1	1008	0x80000002	2001:db8:10:2:1::1/128
10.1.1.1	1007	0x80000001	2001:db8:10:2:10::1/128
10.1.1.1	533	0x80000001	2001:db8:192:168:134::/80
10.1.1.1	523	0x80000001	2001:db8:192:168:41::/80
10.1.1.1	523	0x80000001	2001:db8:10:4:1::1/128
10.1.1.1	523	0x80000001	2001:db8:192:168:104::/80
10.1.1.1	464	0x80000001	2001:db8:10:4:10::1/128
10.1.1.1	454	0x80000001	2001:db8:192:168::/64

Link (Type-8) Link States (Area 11)

ADV Router	Age	Seq#	Link ID	Interface
10.1.1.1	1069	0x80000001	7	Gi0/0/0/0
10.1.10.1	1080	0x80000001	2	Gi0/0/0/0

Intra Area Prefix Link States (Area 11)

ADV Router	Age	Seq#	Link ID	Ref-lstype	Ref-LSID
10.1.10.1	1081	0x80000001	0	0x2001	0

- 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:

```
CE1#show ipv6 route ospf
IPv6 Routing Table - default - 7 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
        B - BGP, R - RIP, I1 - ISIS L1, I2 - ISIS L2
        IA - ISIS interarea, IS - ISIS summary, D - EIGRP, EX - EIGRP external
        ND - Neighbor Discovery
        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:DB8:10:0:1::1/128 [110/2]
    via FE80::4255:39FF:FE2E:C420, GigabitEthernet0/0
OI 2001:DB8:10:0:2::1/128 [110/2]
    via FE80::4255:39FF:FE2E:C420, GigabitEthernet0/0
OI 2001:DB8:10:1:1::1/128 [110/1]
    via FE80::4255:39FF:FE2E:C420, GigabitEthernet0/0
OI 2001:DB8:10:2:1::1/128 [110/2]
    via FE80::4255:39FF:FE2E:C420, GigabitEthernet0/0
OI 2001:DB8:10:2:10::1/128 [110/3]
    via FE80::4255:39FF:FE2E:C420, GigabitEthernet0/0
```

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

```
RP/0/RSP0/CPU0:PE1#show route ipv6 ospf
O    2001:db8:10:0:1::1/128
    [110/1] via fe80::4255:39ff:fe2e:d822, 00:26:58, GigabitEthernet0/0/0/2
O    2001:db8:10:0:2::1/128
    [110/1] via fe80::4255:39ff:fe2f:42da, 00:26:58, GigabitEthernet0/0/0/3
O    2001:db8:10:1:10::1/128
    [110/1] via fe80::eab7:48ff:fe2c:a180, 00:16:45, GigabitEthernet0/0/0/0
O IA 2001:db8:10:2:1::1/128
    [110/1] via fe80::eab7:48ff:fe2c:5801, 00:22:57, GigabitEthernet0/0/0/1
O IA 2001:db8:10:2:10::1/128
```

```
[110/2] via fe80::eab7:48ff:feff:5801, 00:18:32, GigabitEthernet0/0/0/1
■ Verify IPv6 connectivity between 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.

CE1#ping 2001:DB8:10:2:10::1 (IPv6 address of the CE2 router Lo0 interface)

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:10:2:10::1, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 0/0/4 ms
```

Task 2: Influence OSPF DR and BDR Election

In this task, you will influence OSPF DR and BDR election on the LAN. On the LAN segment between CE and PE routers, only the PE router will become DR—the CE router will never become 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.

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 DR.
- Step 2** On the PE router Gi0/0/0/2 and Gi0/0/0/3 interfaces, change the OSPF priority to the value one higher than default.

Activity Verification

You have completed this task when you attain these results:

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

```
RP/0/RSP0/CPU0:PE1#show ospf interface GigabitEthernet 0/0/0/0 | include State
Transmit Delay is 1 sec, State DR, Priority 1, MTU 1500, MaxPktSz 1500
```

- 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:PE1#show ospf interface GigabitEthernet 0/0/0/2 | include State
Transmit Delay is 1 sec, State BDR, Priority 2, MTU 1500, MaxPktSz 1500
RP/0/RSP0/CPU0:PE1#show ospf interface GigabitEthernet 0/0/0/3 | include State
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* DR, it should change when the current DR will go down.

Task 3: Influence OSPF Route Selection by Changing OSPF Link Cost

In this task, you will change the OSPF link cost to influence the OSPF route selection. The OSPF 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.

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:PE1#show route 10.2.1.1
Routing entry for 10.2.1.1/32
  Known via "ospf 1", distance 110, metric 2, type inter area
  Installed Jun  8 01:19:33.142 for 01:51:13
  Routing Descriptor Blocks
    192.168.112.20, from 2.2.2.2, via GigabitEthernet0/0/0/1
    Route metric is 2
  No advertising protos.
RP/0/RSP0/CPU0:PE1#show route 10.2.10.1
Routing entry for 10.2.10.1/32
  Known via "ospf 1", distance 110, metric 3, type inter area
  Installed Jun  8 01:20:08.609 for 01:50:55
  Routing Descriptor Blocks
    192.168.112.20, from 2.2.2.2, 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.

Activity Verification

You have completed this task when you attain these results:

- 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:PE1#show route 10.2.1.1
Routing entry for 10.2.1.1/32
  Known via "ospf 1", distance 110, metric 5, type inter area
  Installed Nov 15 19:43:15.469 for 00:01:47
  Routing Descriptor Blocks
    192.168.11.1, from 2.2.2.2, via GigabitEthernet0/0/0/2
    Route metric is 5
    192.168.12.2, from 2.2.2.2, via GigabitEthernet0/0/0/3
    Route metric is 5
  No advertising protos.
RP/0/RSP0/CPU0:PE1#show route 10.2.10.1
Routing entry for 10.2.10.1/32
  Known via "ospf 1", distance 110, metric 6, type inter area
  Installed Nov 15 19:43:15.469 for 00:02:04
  Routing Descriptor Blocks
    192.168.11.1, from 2.2.2.2, via GigabitEthernet0/0/0/2
    Route metric is 6
    192.168.12.2, from 2.2.2.2, via GigabitEthernet0/0/0/3
    Route metric is 6
  No advertising protos.
```

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

```
PE2#show ip route 10.1.1.1
Routing entry for 10.1.1.1/32
  Known via "ospf 1", distance 110, metric 5, type inter area
  Last update from 192.168.22.2 on GigabitEthernet0/0/3, 00:01:49 ago
  Routing Descriptor Blocks:
    192.168.22.2, from 1.1.1.1, 00:01:49 ago, via GigabitEthernet0/0/3
    Route metric is 5, traffic share count is 1
  * 192.168.21.1, from 1.1.1.1, 00:01:49 ago, via GigabitEthernet0/0/2
    Route metric is 5, traffic share count is 1
PE2#show ip route 10.1.10.1
Routing entry for 10.1.10.1/32
  Known via "ospf 1", distance 110, metric 6, type inter area
```

```

Last update from 192.168.22.2 on GigabitEthernet0/0/3, 00:02:06 ago
Routing Descriptor Blocks:
  192.168.22.2, from 1.1.1.1, 00:02:06 ago, via GigabitEthernet0/0/3
    Route metric is 6, traffic share count is 1
  * 192.168.21.1, from 1.1.1.1, 00:02:06 ago, via GigabitEthernet0/0/2
    Route metric is 6, traffic share count is 1

```

- Trace between 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:PE1#traceroute 10.2.1.1
Type escape sequence to abort.
Tracing the route to 10.2.1.1

```

```

1  192.168.12.2 2 msec
   192.168.11.1 1 msec
   192.168.12.2 0 msec
2  192.168.21.20 0 msec
   192.168.22.20 0 msec  *

```

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

```

PE2#traceroute 10.1.1.1

Type escape sequence to abort.
Tracing the route to 10.1.1.1
VRF info: (vrf in name/id, vrf out name/id)
 1 192.168.21.1 1 msec
   192.168.22.2 1 msec
   192.168.21.1 1 msec
 2 192.168.12.10 40 msec
   192.168.11.10 1 msec
   192.168.12.10 1 msec

```

Task 4: Configure OSPF Authentication

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

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”.

Activity Verification

You have completed this task when you attain these results:

- On the PE router in your pod, verify that 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:PE1#show ospf neighbor 10.10.10.10 | include State
Neighbor priority is 0, State is FULL, 6 state changes

```

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

```

PE2#show ip ospf neighbor 20.20.20.20 | include State
Neighbor priority is 1, State is FULL, 6 state changes

```

- 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:PE1#show ospf interface GigabitEthernet 0/0/0/0
GigabitEthernet0/0/0/0 is up, line protocol is up
Internet Address 192.168.101.10/24, Area 11

```

```

Process ID 1, Router ID 1.1.1.1, Network Type BROADCAST, Cost: 1
Transmit Delay is 1 sec, State DR, Priority 1, MTU 1500, MaxPktSz 1500
Designated Router (ID) 1.1.1.1, Interface address 192.168.101.10
No backup designated router on this network
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
  Hello due in 00:00:08
Index 1/5, flood queue length 0
Next 0(0)/0(0)
Last flood scan length is 12, maximum is 27
Last flood scan time is 0 msec, maximum is 0 msec
LS Ack List: current length 0, high water mark 33
Neighbor Count is 1, Adjacent neighbor count is 1
  Adjacent with neighbor 10.10.10.10
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:

```

PE2#show ip ospf interface GigabitEthernet 0/0/0
GigabitEthernet0/0/0 is up, line protocol is up
  Internet Address 192.168.102.20/24, Area 12, Attached via Network Statement
  Process ID 1, Router ID 2.2.2.2, Network Type BROADCAST, Cost: 1
  Topology-MTID      Cost      Disabled      Shutdown      Topology Name
        0             1          no            no            Base
  Transmit Delay is 1 sec, State BDR, Priority 1
  Designated Router (ID) 20.20.20.20, Interface address 192.168.102.21
  Backup Designated router (ID) 2.2.2.2, Interface address 192.168.102.20
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    oob-resync timeout 40
    Hello due in 00:00:05
  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 14, maximum is 26
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 20.20.20.20 (Designated Router)
  Suppress hello for 0 neighbor(s)
Message digest authentication enabled
  Youngest key id is 1

```

Task 5: Configure 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 CE and PE routers.

Activity Procedure

Complete these steps:

- Step 1** On the CE router, enable the Loopback1x or Loopback1y interface and assign the 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 the OSPF area x00 or y00.
- Step 3** Configure the OSPF virtual link between CE and PE routers.

Activity Verification

You have completed this task when you attain these results:

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

```
CE1#show ip ospf virtual-links
Virtual Link OSPF_VL2 to router 1.1.1.1 is up
  Run as demand circuit
  DoNotAge LSA allowed.
  Transit area 11, 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 0
  First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)
  Last retransmission scan length is 0, maximum is 0
  Last retransmission scan time is 0 msec, maximum is 0 msec
CE1#show ip ospf neighbor
Neighbor ID      Pri   State           Dead Time   Address        Interface
1.1.1.1          0     FULL/-          -           192.168.101.10 OSPF_VL2
1.1.1.1          1     FULL/DR         00:00:31    192.168.101.10 Gi0/0
```

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

```
CE1#show ip ospf database

OSPF Router with ID (10.10.10.10) (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        224      (DNA) 0x80000C15 0x00F389 9
0.0.0.2          0.0.0.2        1141     (DNA) 0x80000C0D 0x00A165 8
1.1.1.1          1.1.1.1        1        (DNA) 0x8000000C 0x001717 4
2.2.2.2          2.2.2.2        2        (DNA) 0x80000007 0x00862B 4
3.3.3.3          3.3.3.3        31       (DNA) 0x80000013 0x00AA2F 4
4.4.4.4          4.4.4.4        31       (DNA) 0x80000007 0x00B6DC 3
5.5.5.5          5.5.5.5        150      (DNA) 0x80000007 0x00FBF9 3
10.7.1.1         10.7.1.1       1721     (DNA) 0x800003CD 0x00B922 3
10.7.10.1        10.7.10.1      383      (DNA) 0x800003C2 0x00C60B 2
10.8.10.1        10.8.10.1      1263     (DNA) 0x800005CF 0x00C0E7 2
10.10.10.10      10.10.10.10    87       (DNA) 0x80000002 0x0023DD 1
10.100.100.100   10.100.100.100 1242     (DNA) 0x800003D4 0x003317 3
20.20.20.20      20.20.20.20    3        (DNA) 0x80000002 0x00E3BD 1
30.30.30.30      30.30.30.30    4        (DNA) 0x80000006 0x009CA1 1

Net Link States (Area 0)

Link ID          ADV Router      Age      Seq#           Checksum
192.168.2.2      0.0.0.2        1141     (DNA) 0x80000B7F 0x00BF7E
192.168.11.1     0.0.0.1        1205     (DNA) 0x80000006 0x009623
192.168.12.2     0.0.0.2        1141     (DNA) 0x80000006 0x008530
192.168.21.1     0.0.0.1        1205     (DNA) 0x80000004 0x005E4F
192.168.22.2     0.0.0.2        1141     (DNA) 0x80000004 0x004D5C
192.168.31.1     0.0.0.1        450      (DNA) 0x80000008 0x001A81
192.168.32.2     0.0.0.2        395      (DNA) 0x80000008 0x00098E
192.168.51.1     0.0.0.1        1205     (DNA) 0x80000005 0x00A7DA
192.168.52.2     0.0.0.2        1405     (DNA) 0x80000005 0x0096E7
192.168.71.70    10.7.1.1       1721     (DNA) 0x800003BE 0x00E278
192.168.82.80    10.100.100.100 1242     (DNA) 0x800003BC 0x0023DB
192.168.107.70   10.7.1.1       1213     (DNA) 0x80000396 0x00A2A1
192.168.108.81   10.8.10.1      1263     (DNA) 0x800003B9 0x005B8A
192.168.112.10   1.1.1.1        1342     (DNA) 0x80000002 0x001B2A
192.168.134.40   4.4.4.4        31       (DNA) 0x80000001 0x003BBA
```

Summary Net Link States (Area 0)

Link ID	ADV Router	Age	Seq#	Checksum
10.1.1.1	1.1.1.1	1342 (DNA)	0x80000003	0x00CD5B
10.1.10.1	1.1.1.1	1342 (DNA)	0x80000002	0x0076A9
10.1.10.1	10.10.10.10	113	0x80000002	0x005D9F
10.2.1.1	2.2.2.2	1283 (DNA)	0x80000002	0x00A57F
10.2.10.1	2.2.2.2	1283 (DNA)	0x80000002	0x004CCE
10.2.10.1	20.20.20.20	9 (DNA)	0x80000001	0x0026AE
10.3.1.1	3.3.3.3	296 (DNA)	0x80000005	0x0075A7
10.3.10.1	3.3.3.3	59 (DNA)	0x80000004	0x001EF5
10.3.10.1	30.30.30.30	86 (DNA)	0x80000002	0x00EABF
10.4.1.1	4.4.4.4	1401 (DNA)	0x80000002	0x0051C9
10.4.10.1	4.4.4.4	1139 (DNA)	0x80000002	0x00F719
10.5.10.1	5.5.5.5	786 (DNA)	0x80000003	0x00CB3F
10.5.10.1	50.50.50.50	4 (DNA)	0x80000001	0x007ADE
11.0.0.1	10.10.10.10	114	0x80000002	0x00CA3C
12.0.0.1	20.20.20.20	9 (DNA)	0x80000001	0x00924C
13.0.0.1	30.30.30.30	96 (DNA)	0x80000001	0x00585D
192.168.101.0	1.1.1.1	603 (DNA)	0x80000008	0x005A08
192.168.101.0	10.10.10.10	114	0x80000002	0x0057EC
192.168.102.0	2.2.2.2	1283 (DNA)	0x80000002	0x003D26
192.168.102.0	20.20.20.20	9 (DNA)	0x80000001	0x0021FA
192.168.103.0	3.3.3.3	59 (DNA)	0x80000006	0x000C4E
192.168.103.0	30.30.30.30	86 (DNA)	0x80000002	0x00E60B
192.168.104.0	4.4.4.4	1139 (DNA)	0x80000002	0x00EA6E
192.168.105.0	5.5.5.5	1162 (DNA)	0x80000007	0x00B797
192.168.105.0	50.50.50.50	4 (DNA)	0x80000001	0x007828

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	952 (DNA)	0x800000E2	0x00E35F	0
1.0.0.0	0.0.0.2	1141 (DNA)	0x800000E2	0x00E65A	0
1.0.0.0	10.7.1.1	1213 (DNA)	0x800000E2	0x0082A7	0
1.0.0.0	10.100.100.100	1242 (DNA)	0x800000E2	0x009371	0
1.0.0.4	10.100.100.100	1242 (DNA)	0x800000E7	0x002816	4
1.0.0.8	0.0.0.1	952 (DNA)	0x800000E7	0x00CBFA	8
1.0.0.9	10.7.1.1	206 (DNA)	0x800000E8	0x004C18	9
1.0.0.11	0.0.0.2	1141 (DNA)	0x800000E9	0x00641A	11
1.0.0.12	0.0.0.1	952 (DNA)	0x800000E9	0x004A2C	12
1.0.0.19	0.0.0.2	1141 (DNA)	0x800000E7	0x00B74E	19

Router Link States (Area 11)

Link ID	ADV Router	Age	Seq#	Checksum	Link count
1.1.1.1	1.1.1.1	887	0x80000005	0x00FC77	1
10.10.10.10	10.10.10.10	110	0x80000008	0x00F606	2

Net Link States (Area 11)

Link ID	ADV Router	Age	Seq#	Checksum
192.168.101.10	1.1.1.1	887	0x80000001	0x002809

Summary Net Link States (Area 11)

Link ID	ADV Router	Age	Seq#	Checksum
10.0.1.1	1.1.1.1	1453	0x80000003	0x00E345
10.0.2.1	1.1.1.1	729	0x80000006	0x00E63C
10.1.1.1	1.1.1.1	1453	0x80000003	0x00CD5B
10.2.1.1	1.1.1.1	729	0x80000008	0x00DF3F
10.2.10.1	1.1.1.1	729	0x80000003	0x009089
10.3.1.1	1.1.1.1	1454	0x80000003	0x00C95B
10.3.10.1	1.1.1.1	174	0x80000001	0x0074A8
10.4.1.1	1.1.1.1	140	0x80000001	0x0026F5
10.4.10.1	1.1.1.1	140	0x80000001	0x00CC45
10.5.10.1	1.1.1.1	895	0x80000003	0x0058C0
10.7.1.1	1.1.1.1	1454	0x80000004	0x00F125
10.7.10.1	1.1.1.1	1454	0x80000003	0x009A73
10.8.1.1	1.1.1.1	730	0x80000006	0x00F51C

10.8.10.1	1.1.1.1	730	0x80000006	0x009C6B
11.0.0.1	1.1.1.1	112	0x80000001	0x00E545
11.0.0.1	10.10.10.10	127	0x80000001	0x00CC3B
12.0.0.1	1.1.1.1	97	0x80000001	0x000125
13.0.0.1	1.1.1.1	168	0x80000001	0x00DF47
192.168.2.0	1.1.1.1	1454	0x80000003	0x00BD0B
192.168.11.0	1.1.1.1	1454	0x80000005	0x00427D
192.168.12.0	1.1.1.1	760	0x80000006	0x004972
192.168.21.0	1.1.1.1	751	0x80000008	0x00EBC3
192.168.22.0	1.1.1.1	731	0x80000005	0x00E6CA
192.168.31.0	1.1.1.1	1455	0x80000003	0x007339
192.168.32.0	1.1.1.1	558	0x80000004	0x008423
192.168.41.0	1.1.1.1	1455	0x80000003	0x001987
192.168.42.0	1.1.1.1	731	0x80000006	0x000894
192.168.51.0	1.1.1.1	1455	0x80000003	0x009602
192.168.52.0	1.1.1.1	761	0x80000004	0x00A7EB
192.168.61.0	1.1.1.1	337	0x80000001	0x00404E
192.168.71.0	1.1.1.1	1455	0x80000003	0x001467
192.168.82.0	1.1.1.1	731	0x80000006	0x00A8C2
192.168.102.0	1.1.1.1	731	0x80000003	0x0081E0
192.168.103.0	1.1.1.1	219	0x80000001	0x0066FE
192.168.104.0	1.1.1.1	142	0x80000001	0x00BF9A
192.168.105.0	1.1.1.1	1455	0x80000003	0x004C15
192.168.107.0	1.1.1.1	1456	0x80000003	0x0090C5
192.168.108.0	1.1.1.1	732	0x80000006	0x0093BC
192.168.112.0	1.1.1.1	695	0x8000000A	0x003715
192.168.134.0	1.1.1.1	141	0x80000005	0x0062D6
192.168.156.0	1.1.1.1	192	0x80000005	0x006FB3

Router Link States (Area 100)

Link ID	ADV Router	Age	Seq#	Checksum	Link count
10.10.10.10	10.10.10.10	129	0x80000002	0x00587B	1

Summary Net Link States (Area 100)

Link ID	ADV Router	Age	Seq#	Checksum
10.0.1.1	10.10.10.10	110	0x80000001	0x00E223
10.0.2.1	10.10.10.10	110	0x80000001	0x00EB17
10.1.1.1	10.10.10.10	110	0x80000001	0x00CC39
10.1.10.1	10.10.10.10	130	0x80000001	0x005F9E
10.2.1.1	10.10.10.10	110	0x80000001	0x00E818
10.2.10.1	10.10.10.10	110	0x80000001	0x008F67
10.3.1.1	10.10.10.10	110	0x80000001	0x00C839
10.3.10.1	10.10.10.10	110	0x80000001	0x006F88
10.4.1.1	10.10.10.10	110	0x80000001	0x0021D5
10.4.10.1	10.10.10.10	110	0x80000001	0x00C725
10.5.10.1	10.10.10.10	110	0x80000001	0x00579E
10.7.1.1	10.10.10.10	110	0x80000001	0x00F202
10.7.10.1	10.10.10.10	110	0x80000001	0x009951
10.8.1.1	10.10.10.10	110	0x80000001	0x00FAF6
10.8.10.1	10.10.10.10	111	0x80000001	0x00A146
12.0.0.1	10.10.10.10	95	0x80000001	0x00FB05
13.0.0.1	10.10.10.10	111	0x80000001	0x00DA27
192.168.2.0	10.10.10.10	111	0x80000001	0x00BCE8
192.168.11.0	10.10.10.10	111	0x80000001	0x004559
192.168.12.0	10.10.10.10	111	0x80000001	0x004E4D
192.168.21.0	10.10.10.10	111	0x80000001	0x00F49C
192.168.22.0	10.10.10.10	112	0x80000001	0x00E9A6
192.168.31.0	10.10.10.10	112	0x80000001	0x007217
192.168.32.0	10.10.10.10	112	0x80000001	0x0085FF
192.168.41.0	10.10.10.10	112	0x80000001	0x001865
192.168.42.0	10.10.10.10	112	0x80000001	0x000D6F
192.168.51.0	10.10.10.10	112	0x80000001	0x0095DF
192.168.52.0	10.10.10.10	112	0x80000001	0x00A8C8
192.168.61.0	10.10.10.10	112	0x80000001	0x003B2E
192.168.71.0	10.10.10.10	112	0x80000001	0x001345
192.168.82.0	10.10.10.10	112	0x80000001	0x00AD9D

192.168.101.0	10.10.10.10	132	0x80000001	0x0059EB
192.168.102.0	10.10.10.10	112	0x80000001	0x0080BE
192.168.103.0	10.10.10.10	112	0x80000001	0x0061DE
192.168.104.0	10.10.10.10	112	0x80000001	0x00BA7A
192.168.105.0	10.10.10.10	112	0x80000001	0x004BF2
192.168.107.0	10.10.10.10	112	0x80000001	0x008FA3
192.168.108.0	10.10.10.10	113	0x80000001	0x009897
192.168.112.0	10.10.10.10	114	0x80000001	0x0044EB
192.168.134.0	10.10.10.10	114	0x80000001	0x0065B2
192.168.156.0	10.10.10.10	114	0x80000001	0x00728F

- 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:PE1#ping 11.0.0.1
```

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

```
Sending 5, 100-byte ICMP Echos to 11.0.0.1, timeout is 2 seconds:
```

```
!!!!
```

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

Lab 2-2: Implement OSPF Special Area Types

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

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

Cisco IOS XR Commands

Command	Description
<code>[no] virtual-link IP-address</code>	(Router) Enables or disables OSPF virtual link
<code>area area</code>	(Router) Creates OSPF area on the router
<code>commit</code>	Commits changes to the running configuration
<code>configure terminal</code>	Enters configuration mode
<code>default-cost cost</code>	(Router) Changes cost of the default route sent by ABR into stub area
<code>ping dest_IP source [source_IP]</code>	Verifies connectivity between source IP and destination IP
<code>range prefix/subnet-mask</code>	(Router) Summarizes prefixes from nonbackbone OSPF area to the backbone OSPF area

<code>router ospf <i>process-ID</i></code>	Creates an OSPF process
<code>show ospf database</code>	Displays the content of the OSPF database
<code>show route</code>	Displays the current routes in the routing table
<code>stub nssa [no-summary]</code>	(Router) Changes nonbackbone OSPF area into stub area or NSSA. Use the no-summary keyword to create totally stubby area or totally NSSA

Task 1: Enable OSPF Summarization

In this task, you will enable OSPF summarization.

Activity Procedure

Complete these steps:

- Step 1** On the CE and PE routers, remove the OSPF virtual link 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 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 intra-area OSPF route and Loopback1x or Loopback1y network is shown as external OSPF route.

```
RP/0/RSP0/CPU0:PE1#show route ospf
O    10.0.1.1/32 [110/2] via 192.168.11.1, 5d21h, GigabitEthernet0/0/0/2
O    10.0.2.1/32 [110/3] via 192.168.11.1, 5d18h, GigabitEthernet0/0/0/2
O    10.1.10.1/32 [110/2] via 192.168.101.11, 5d17h, GigabitEthernet0/0/0/0
O IA 10.2.1.1/32 [110/4] via 192.168.11.1, 5d18h, GigabitEthernet0/0/0/2
O IA 10.2.10.1/32 [110/5] via 192.168.11.1, 00:13:38, GigabitEthernet0/0/0/2
O E2 11.0.0.0/24 [110/20] via 192.168.101.11, 00:06:36, GigabitEthernet0/0/0/0
O E2 12.0.0.0/24 [110/20] via 192.168.11.1, 00:04:44, GigabitEthernet0/0/0/2
< text omitted >
```

- Step 3** On the PE router, summarize the intra-area OSPF (CE Loopback0 interface network) route to the subnet mask 24.
- Step 4** On the CE router, summarize the external OSPF (CE Loopback1x or Loopback1y interface network) route to the subnet mask 8.

Activity Verification

You have completed this task when you attain these results:

- 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 ABR in the area 1x or 1y (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, taken from the PE1 router:

```
RP/0/RSP0/CPU0:PE1#show route ospf
O    10.0.1.1/32 [110/2] via 192.168.11.1, 5d21h, GigabitEthernet0/0/0/2
O    10.0.2.1/32 [110/3] via 192.168.11.1, 5d18h, GigabitEthernet0/0/0/2
O IA 10.1.10.0/24 [254/0] via 0.0.0.0, 00:02:26, Null0
O    10.1.10.1/32 [110/2] via 192.168.101.11, 5d17h, GigabitEthernet0/0/0/0
O IA 10.2.1.1/32 [110/4] via 192.168.11.1, 5d18h, GigabitEthernet0/0/0/2
O IA 10.2.10.0/24 [110/5] via 192.168.11.1, 00:02:31, GigabitEthernet0/0/0/2
O E2 11.0.0.0/8 [110/20] via 192.168.101.11, 00:04:36, GigabitEthernet0/0/0/0
O E2 12.0.0.0/8 [110/20] via 192.168.11.1, 00:04:05, GigabitEthernet0/0/0/2
< text omitted >
```

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

```
CE1#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, + - replicated route

```

Gateway of last resort is not set

```

10.0.0.0/8 is variably subnetted, 6 subnets, 2 masks
O IA 10.0.1.1/32 [110/3] via 192.168.101.10, 00:25:08, GigabitEthernet0/0
O IA 10.0.2.1/32 [110/4] via 192.168.101.10, 00:25:08, GigabitEthernet0/0
O IA 10.1.1.1/32 [110/2] via 192.168.101.10, 00:25:08, GigabitEthernet0/0
O IA 10.2.1.1/32 [110/5] via 192.168.101.10, 00:25:08, GigabitEthernet0/0
O IA 10.2.10.0/24 [110/6] via 192.168.101.10, 00:09:43, GigabitEthernet0/0
11.0.0.0/8 is variably subnetted, 3 subnets, 3 masks
O 11.0.0.0/8 is a summary, 00:11:48, Null0
O E2 12.0.0.0/8 [110/20] via 192.168.101.10, 00:11:11, GigabitEthernet0/0
< text omitted >

```

Task 2: Configure 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.

Activity Procedure

Complete these steps:

Step 1 On the CE router, examine the IP routing table, and you should see intra-area, interarea, and external OSPF routes. Output should be similar to the following, taken from the CE1 router:

```

CE1#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, + - replicated route

```

Gateway of last resort is not set

```

10.0.0.0/8 is variably subnetted, 6 subnets, 2 masks
O IA 10.0.1.1/32 [110/3] via 192.168.101.10, 01:04:29, GigabitEthernet0/0
O IA 10.0.2.1/32 [110/4] via 192.168.101.10, 01:04:29, GigabitEthernet0/0
O IA 10.1.1.1/32 [110/2] via 192.168.101.10, 01:04:29, GigabitEthernet0/0
C 10.1.10.1/32 is directly connected, Loopback0
O IA 10.2.1.1/32 [110/5] via 192.168.101.10, 01:04:29, GigabitEthernet0/0
O IA 10.2.10.0/24 [110/6] via 192.168.101.10, 00:49:04, GigabitEthernet0/0
11.0.0.0/8 is variably subnetted, 3 subnets, 3 masks
O 11.0.0.0/8 is a summary, 00:51:10, Null0
C 11.0.0.0/24 is directly connected, Loopback11
L 11.0.0.1/32 is directly connected, Loopback11
O E2 12.0.0.0/8 [110/20] via 192.168.101.10, 00:50:33, GigabitEthernet0/0
O IA 192.168.1.0/24 [110/3] via 192.168.101.10, 01:04:30, GigabitEthernet0/0
O IA 192.168.2.0/24 [110/4] via 192.168.101.10, 01:04:30, GigabitEthernet0/0
O IA 192.168.11.0/24 [110/2] via 192.168.101.10, 01:04:30, GigabitEthernet0/0
O IA 192.168.12.0/24 [110/4] via 192.168.101.10, 01:04:30, GigabitEthernet0/0
O IA 192.168.21.0/24 [110/5] via 192.168.101.10, 01:04:30, GigabitEthernet0/0
O IA 192.168.22.0/24 [110/4] via 192.168.101.10, 01:04:30, GigabitEthernet0/0
O IA 192.168.31.0/24 [110/3] via 192.168.101.10, 01:04:30, GigabitEthernet0/0
O IA 192.168.32.0/24 [110/6] via 192.168.101.10, 01:04:30, GigabitEthernet0/0
O IA 192.168.41.0/24 [110/5] via 192.168.101.10, 01:04:30, GigabitEthernet0/0
O IA 192.168.42.0/24 [110/4] via 192.168.101.10, 01:04:30, GigabitEthernet0/0

```

```

O IA 192.168.51.0/24 [110/3] via 192.168.101.10, 01:04:30, GigabitEthernet0/0
O IA 192.168.52.0/24 [110/6] via 192.168.101.10, 01:04:30, GigabitEthernet0/0
O IA 192.168.61.0/24 [110/5] via 192.168.101.10, 01:04:31, GigabitEthernet0/0
O IA 192.168.62.0/24 [110/4] via 192.168.101.10, 01:04:31, GigabitEthernet0/0
O IA 192.168.71.0/24 [110/3] via 192.168.101.10, 01:04:31, GigabitEthernet0/0
O IA 192.168.72.0/24 [110/6] via 192.168.101.10, 01:04:31, GigabitEthernet0/0
O IA 192.168.81.0/24 [110/5] via 192.168.101.10, 01:04:31, GigabitEthernet0/0
O IA 192.168.82.0/24 [110/4] via 192.168.101.10, 01:04:31, GigabitEthernet0/0
192.168.101.0/24 is variably subnetted, 2 subnets, 2 masks
C 192.168.101.0/24 is directly connected, GigabitEthernet0/0
L 192.168.101.11/32 is directly connected, GigabitEthernet0/0
O IA 192.168.102.0/24
[110/5] via 192.168.101.10, 01:04:32, GigabitEthernet0/0
O IA 192.168.112.0/24
[110/11] via 192.168.101.10, 01:04:32, GigabitEthernet0/0

```

Step 2 Configure an OSPF area between the CE and PE routers in your pod to be stub.

Step 3 On the CE router, examine the IP routing table, and you should see 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, taken from the CE1 router:

```

CE1#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, + - replicated route

```

Gateway of last resort is 192.168.101.10 to network 0.0.0.0

```

O*IA 0.0.0.0/0 [110/2] via 192.168.101.10, 00:00:02, GigabitEthernet0/0
10.0.0.0/8 is variably subnetted, 6 subnets, 2 masks
O IA 10.0.1.1/32 [110/3] via 192.168.101.10, 00:00:02, GigabitEthernet0/0
O IA 10.0.2.1/32 [110/4] via 192.168.101.10, 00:00:02, GigabitEthernet0/0
O IA 10.1.1.1/32 [110/2] via 192.168.101.10, 00:00:02, GigabitEthernet0/0
C 10.1.10.1/32 is directly connected, Loopback0
O IA 10.2.1.1/32 [110/5] via 192.168.101.10, 00:00:02, GigabitEthernet0/0
O IA 10.2.10.0/24 [110/6] via 192.168.101.10, 00:00:02, GigabitEthernet0/0
11.0.0.0/8 is variably subnetted, 3 subnets, 3 masks
O 11.0.0.0/8 is a summary, 00:56:00, Null0
C 11.0.0.0/24 is directly connected, Loopback11
L 11.0.0.1/32 is directly connected, Loopback11
O IA 192.168.1.0/24 [110/3] via 192.168.101.10, 00:00:02, GigabitEthernet0/0
O IA 192.168.2.0/24 [110/4] via 192.168.101.10, 00:00:04, GigabitEthernet0/0
O IA 192.168.11.0/24 [110/2] via 192.168.101.10, 00:00:04, GigabitEthernet0/0
O IA 192.168.12.0/24 [110/4] via 192.168.101.10, 00:00:04, GigabitEthernet0/0
O IA 192.168.21.0/24 [110/5] via 192.168.101.10, 00:00:04, GigabitEthernet0/0
O IA 192.168.22.0/24 [110/4] via 192.168.101.10, 00:00:04, GigabitEthernet0/0
O IA 192.168.31.0/24 [110/3] via 192.168.101.10, 00:00:04, GigabitEthernet0/0
O IA 192.168.32.0/24 [110/6] via 192.168.101.10, 00:00:04, GigabitEthernet0/0
O IA 192.168.41.0/24 [110/5] via 192.168.101.10, 00:00:04, GigabitEthernet0/0
O IA 192.168.42.0/24 [110/4] via 192.168.101.10, 00:00:04, GigabitEthernet0/0
O IA 192.168.51.0/24 [110/3] via 192.168.101.10, 00:00:04, GigabitEthernet0/0
O IA 192.168.52.0/24 [110/6] via 192.168.101.10, 00:00:04, GigabitEthernet0/0
O IA 192.168.61.0/24 [110/5] via 192.168.101.10, 00:00:04, GigabitEthernet0/0
O IA 192.168.62.0/24 [110/4] via 192.168.101.10, 00:00:04, GigabitEthernet0/0
O IA 192.168.71.0/24 [110/3] via 192.168.101.10, 00:00:05, GigabitEthernet0/0
O IA 192.168.72.0/24 [110/6] via 192.168.101.10, 00:00:05, GigabitEthernet0/0
O IA 192.168.81.0/24 [110/5] via 192.168.101.10, 00:00:05, GigabitEthernet0/0
O IA 192.168.82.0/24 [110/4] via 192.168.101.10, 00:00:05, GigabitEthernet0/0
192.168.101.0/24 is variably subnetted, 2 subnets, 2 masks
C 192.168.101.0/24 is directly connected, GigabitEthernet0/0
L 192.168.101.11/32 is directly connected, GigabitEthernet0/0

```

```
O IA 192.168.102.0/24
    [110/5] via 192.168.101.10, 00:00:05, GigabitEthernet0/0
O IA 192.168.112.0/24
    [110/11] via 192.168.101.10, 00:00:05, 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 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, taken from the CE1 router:

```
CE1#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, + - replicated route
```

Gateway of last resort is 192.168.101.10 to network 0.0.0.0

```
O*IA 0.0.0.0/0 [110/2] via 192.168.101.10, 00:01:11, GigabitEthernet0/0
      10.0.0.0/32 is subnetted, 1 subnets
C      10.1.10.1 is directly connected, Loopback0
      11.0.0.0/8 is variably subnetted, 3 subnets, 3 masks
O      11.0.0.0/8 is a summary, 01:01:45, Null0
C      11.0.0.0/24 is directly connected, Loopback11
L      11.0.0.1/32 is directly connected, Loopback11
      192.168.101.0/24 is variably subnetted, 2 subnets, 2 masks
C      192.168.101.0/24 is directly connected, GigabitEthernet0/0
L      192.168.101.11/32 is directly connected, GigabitEthernet0/0
```

Step 6 On the PE router, adjust the OSPF cost of 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, taken from the CE1 router:

```
CE1#show ip ospf database summary 0.0.0.0

        OSPF Router with ID (10.10.10.10) (Process ID 1)

        Summary Net Link States (Area 11)

Routing Bit Set on this LSA in topology Base with MTID 0
LS age: 615
Options: (No TOS-capability, DC, Upward)
LS Type: Summary Links(Network)
Link State ID: 0.0.0.0 (summary Network Number)
Advertising Router: 1.1.1.1
LS Seq Number: 80000003
Checksum: 0xE945
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, taken from the CE1 router:

```
CE1#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 192.168.101.10 on GigabitEthernet0/0, 00:11:56 ago
```

Routing Descriptor Blocks:

* 192.168.101.10, from 1.1.1.1, 00:11:56 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 stub, there is no IP reachability to the Loopback1x or Loopback1y (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.

Activity Verification

You have completed this task when you attain these results:

- On the PE router in your pod, verify the OSPF database. There should be type 7 LSA, originated by the CE router. Type 7 LSA describes the CE Loopback1x or Loopback1y network (x or y is your pod number). Output should be similar to the following, taken from the PE1 router:

```
RP/0/RSP0/CPU0:PE1#show ospf database nssa-external
      OSPF Router with ID (1.1.1.1) (Process ID 1)

      Type-7 AS External Link States (Area 11)

Routing Bit Set on this LSA
LS age: 115
Options: (No TOS-capability, Type 7/5 translation, DC)
LS Type: AS External Link
Link State ID: 11.0.0.0 (External Network Number)
Advertising Router: 10.10.10.10
LS Seq Number: 80000001
Checksum: 0xalab
Length: 36
Network Mask: /8
      Metric Type: 2 (Larger than any link state path)
      TOS: 0
      Metric: 20
      Forward Address: 10.1.10.1
      External Route Tag: 0
```

- The PE router translates type 7 LSA into type 5 LSA. On the PE router in your pod, verify the OSPF database. There should be type 5 LSA originated by the PE router. Type 5 LSA describes the CE Loopback1x or Loopback1y network (x or y is your pod number). Outputs should be similar to the following, taken from the PE1 router:

```
RP/0/RSP0/CPU0:PE1#show ospf database external 11.0.0.0
      OSPF Router with ID (1.1.1.1) (Process ID 1)

      Type-5 AS External Link States

LS age: 368
Options: (No TOS-capability, DC)
LS Type: AS External Link
Link State ID: 11.0.0.0 (External Network Number)
Advertising Router: 1.1.1.1
LS Seq Number: 80000001
Checksum: 0x4536
Length: 36
Network Mask: /8
      Metric Type: 2 (Larger than any link state path)
      TOS: 0
      Metric: 20
      Forward Address: 10.1.10.1
      External Route Tag: 0
```

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

```
CE1#ping 12.0.0.1 (Lo12 IP address from CE2 router)
```

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

```
Sending 5, 100-byte ICMP Echos to 12.0.0.1, timeout is 2 seconds:
```

```
!!!!
```

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

Lab 3-1: Implement Integrated IS-IS Routing

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

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 in 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.



NET Addressing

This section includes a table with NET addresses used in the lab.

Pod NET Addresses

Replace the x or y with your pod number to get router NET addresses.

Router	NET address
CEx, CEy	49.000x.0100.0x01.0001.00, 49.000y.0100.0y01.0001.00
PEx, PEy	49.000x.0100.0x00.1001.00, 49.000y.0100.0y00.1001.00

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

Cisco IOS XR Commands

Command	Description
address-family <i>ipv4 ipv6</i> unicast	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 IS-IS routing process to establish selected IS-IS circuit-type on the interface
commit	Commit 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
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 NET address to the IS-IS router
ping <i>dest_IP</i> <i>source source_IP</i>	Verifies connectivity between source IP and destination IP (IPv4 and IPv6)
router isis <i>process-ID</i>	Creates an IS-IS process
show ipv4 interface brief	Displays interface status and IPv4 addresses 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 related to the IS-IS topology for a specific router
show route	Displays the current routes in the routing table
show running-config	Displays running configuration
single-topology	Enables IS-IS to support single topology. Used in the IPv6 IS-IS address family
summary-prefix <i>network/subnet_mask</i> <i>level-1 level-1-2 level-2</i>	Summarizes IS-IS route to the selected level. Default is summarizing into level 2

Task 1: Enable 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.

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 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 be first Gigabit Ethernet and Loopback interfaces up and running with assigned IP addresses.

```
CE1#show ip interface brief | include GigabitEthernet0/0|Loopback0
GigabitEthernet0/0      192.168.101.11 YES manual up      up
Loopback0               10.1.10.1      YES manual 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:PE1#show ipv4 interface brief | include Up
Sat May 20 05:05:30.676 UTC
Loopback0                10.1.1.1        Up              Up
MgmtEth0/RSP0/CPU0/0     10.10.10.17     Up              Up
GigabitEthernet0/0/0/0    192.168.101.10  Up              Up
GigabitEthernet0/0/0/1    192.168.112.10  Up              Up
GigabitEthernet0/0/0/2    192.168.11.10   Up              Up
GigabitEthernet0/0/0/3    192.168.12.10   Up              Up
```

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

Successful ping from PE to CE router.

```
RP/0/RSP0/CPU0:PE1#ping 192.168.101.11
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.101.11, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

Successful ping between PE routers in your team.

```
RP/0/RSP0/CPU0:PE1#ping 192.168.112.20
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.112.20, 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.

Activity Verification

You have completed this task when you attain these results:

- 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 CE and PE routers. In the IS-IS database, there should be three LSPs. Outputs should be similar to the following:

```
CE1#show isis neighbors
System Id      Type Interface  IP Address      State Holdtime Circuit Id
PE1            L1  Gi0/0        192.168.101.10  UP      21         CE1.01

CE1#show isis topology
IS-IS TID 0 paths to level-1 routers
System Id      Metric  Next-Hop      Interface      SNPA
PE1            10     PE1           Gi0/0          4055.392e.c420
CE1            --

CE1#show isis database
IS-IS Level-1 Link State Database:
LSPID          LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
PE1.00-00      0x000001CD  0xEB68        671           1/0/0
CE1.00-00      * 0x000001C8  0x2E21        1170          0/0/0
CE1.01-00      * 0x000001B8  0xADD6        483           0/0/0
```

- On the PE router in your pod, verify IS-IS neighbors, IS-IS topology, and IS-IS database. 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 CE and PE routers in the Level 1 topology and 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. Outputs should be similar to the following:

```
RP/0/RSP0/CPU0:PE1#show isis neighbors
IS-IS 1 neighbors:
System Id      Interface      SNPA              State Holdtime Type IETF-NSF
CE1            Gi0/0/0/0     e8b7.482c.a180   Up      9      L1    Capable
PE2            Gi0/0/0/1     e8b7.48fb.5801   Up      7      L2    Capable

Total neighbor count: 2
RP/0/RSP0/CPU0:PE1#show isis topology
IS-IS 1 paths to IPv4 Unicast (Level-1) routers
System Id      Metric  Next-Hop      Interface      SNPA
PE1            --
CE1            10     CE1           Gi0/0/0/0     e8b7.482c.a180
```

```
IS-IS 1 paths to IPv4 Unicast (Level-2) routers
System Id      Metric  Next-Hop      Interface      SNPA
PE1            --
CE1            **
PE2            10      PE2            Gi0/0/0/1      e8b7.48fb.5801
```

```
RP/0/RSP0/CPU0:PE1#show isis database
```

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

LSPID	LSP Seq Num	LSP Checksum	LSP Holdtime	ATT/P/OL
PE1.00-00	* 0x000001ce	0xe969	1180	1/0/0
CE1.00-00	0x000001c8	0x2e21	975	0/0/0
CE1.01-00	0x000001b9	0xabd7	1028	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
PE1.00-00	* 0x000001d3	0x0265	479	0/0/0
CE1.00-00	0x0000000c	0x55e6	233	0/0/0
CE1.01-00	0x00000009	0xa11a	115	0/0/0
PE2.00-00	0x000001ce	0x538f	1157	0/0/0
PE2.02-00	0x00000008	0x29b5	0 (296)	0/0/0
PE2.03-00	0x000001b6	0x7377	478	0/0/0

```
Total Level-2 LSP count: 6      Local Level-2 LSP count: 1
```

- 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:

```
CE1#show ip route isis
< text omitted >
```

```
Gateway of last resort is 192.168.101.10 to network 0.0.0.0
```

```
i*L1  0.0.0.0/0 [115/10] via 192.168.101.10, GigabitEthernet0/0
      10.0.0.0/32 is subnetted, 2 subnets
i L1   10.1.1.1 [115/20] via 192.168.101.10, GigabitEthernet0/0
```

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

```
RP/0/RSP0/CPU0:PE1#show route isis
```

```
i su 10.1.10.0/24 [115/20] via 0.0.0.0, 00:16:11, Null0
i L1 10.1.10.1/32 [115/20] via 192.168.101.11, 00:03:22,
GigabitEthernet0/0/0/0
i L2 10.2.1.0/24 [115/20] via 192.168.112.20, 00:22:52, GigabitEthernet0/0/0/1
i L2 10.2.10.0/24 [115/30] via 192.168.112.20, 00:03:16,
GigabitEthernet0/0/0/1
i L2 192.168.102.0/24 [115/20] via 192.168.112.20, 00:03:12,
GigabitEthernet0/0/0/1
```

- Verify IPv4 connectivity between 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.

```
CE1#ping 10.2.10.1 (IP address of the CE2 router Loopback0 interface)
```

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

```
Sending 5, 100-byte ICMP Echos to 10.2.10.1, timeout is 2 seconds:
```

```
!!!!
```

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

Activity Verification

You have completed this task when you attain these results:

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

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

```
RP/0/RSP0/CPU0:PE1#show route isis | include su
i su 10.1.10.0/24 [115/20] via 0.0.0.0, 00:06:39, Null0
```

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

```
PE2#show ip route isis | include i su
i su 10.2.10.0/24 [115/20] via 0.0.0.0, 00:05:57, Null0
```

- 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:PE1#show route isis | include 10.0/24
i su 10.1.10.0/24 [115/20] via 0.0.0.0, 00:10:49, Null0
i L2 10.2.10.0/24 [115/30] via 192.168.112.20, 00:09:08, Gi0/0/0/1
```

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

```
PE2#show ip route isis
< text omitted >
10.0.0.0/8 is variably subnetted, 6 subnets, 2 masks
i L2 10.1.1.1/32
[115/20] via 192.168.112.10, 02:32:27, GigabitEthernet0/0/1
i L2 10.1.10.0/24
[115/30] via 192.168.112.10, 00:11:57, GigabitEthernet0/0/1
i su 10.2.10.0/24 [115/20] via 0.0.0.0, 00:10:18, Null0
i L1 10.2.10.1/32
[115/20] via 192.168.102.21, 00:48:30, GigabitEthernet0/0/0
i L2 192.168.101.0/24
[115/20] via 192.168.112.10, 02:32:27, GigabitEthernet0/0/1
```

- Verify IPv4 connectivity between 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 still be successful.

CE1#ping 10.2.10.1 (IP address of the CE2 router Loopback0 interface)

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 10.2.10.1, timeout is 2 seconds:

!!!!

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

Task 3: Enable 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 routers Loopback interfaces. Find 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 (pod number may be different):

```
ipv6 unicast-routing
!
interface Loopback0
  ipv6 address 2001:DB8:10:1:10::1/128
  ipv6 enable
!
interface GigabitEthernet0/0
  ipv6 enable
```

If IPv6 is not preconfigured on the PE router, enter commands similar to these (pod number may be different):

PE1 (Cisco IOS XR):

```
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
```

PE2 (Cisco IOS XE):

```
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.

Activity Verification

You have completed this task when you attain these results:

- 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:

```
CE1#show ipv6 route isis
< text omitted >
i1  ::/0 [115/10]
    via FE80::4255:39FF:FE2E:C420, GigabitEthernet0/0
i1  2001:DB8:10:1:1::1/128 [115/20]
    via FE80::4255:39FF:FE2E:C420, GigabitEthernet0/0
```

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

```
RP/0/RSP0/CPU0:PE1#show route ipv6 isis
i L1 2001:db8:10:1:10::1/128
    [115/20] via fe80::eab7:48ff:fe2c:a180, 00:00:30, GigabitEthernet0/0/0/0
i L2 2001:db8:10:2:1::1/128
    [115/20] via fe80::eab7:48ff:fe2c:a180, 00:00:24, GigabitEthernet0/0/0/1
i L2 2001:db8:10:2:10::1/128
    [115/30] via fe80::eab7:48ff:fe2c:a180, 00:00:02, GigabitEthernet0/0/0/1
```

- Verify IPv6 connectivity between 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.

```
CE1#ping 2001:DB8:10:2:10::1 (IPv6 address of the CE2 router Loopback0
interface)
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:10:2:10::1, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 0/1/8 ms
```

Lab 4-1: Implement Basic BGP Routing

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

Activity Objective

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


Note Students from two different pods are working in 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.



NET Addressing

This section includes a table with NET addresses used in the lab.

Pod and Backbone NET Addresses

Replace the x or y with your pod number to get the NET addresses for routers.

Router	NET address
P1	49.0000.0100.0000.1001.00
P2	49.0000.0100.0000.2001.00
CEx, CEy	49.000x.0100.0x01.0001.00, 49.000y.0100.0y01.0001.00
PEx, PEy	49.000x.0100.0x00.1001.00, 49.000y.0100.0y00.1001.00

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 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
<code>[no] shutdown</code>	Enables or disables interface on the router
<code>cdp enable</code>	Enables Cisco Discovery Protocol on the interface
<code>configure terminal</code>	Enters configuration mode
<code>interface interface</code>	Enters interface configuration mode
<code>ip address ip_address subnet_mask</code>	Sets a primary or secondary IPv4 address for an interface and the subnet mask
<code>ip ipv6 router isis</code>	Enables IS-IS routing to the interface
<code>ipv6 enable</code>	Enables IPv6 support on the interface
<code>ipv6 unicast-routing</code>	Enables IPv6 unicast routing (global)
<code>isis circuit-type level-1 level-1-2 level-2-only</code>	Enables IS-IS routing process to establish selected IS-IS circuit-type on the interface
<code>neighbor IP-address next-hop-self</code>	Enables internal BGP router to send BGP routes with router BGP own IP address
<code>neighbor IP-address password password</code>	Enables authentication for BGP neighbor
<code>neighbor IP-address remote-as AS-number</code>	Configures BGP neighbor
<code>neighbor IP-address update-source interface</code>	Enables BGP router to send BGP packets from source IP address
<code>network network mask subnet-mask</code>	Enables BGP router to include network from routing table into BGP routing
<code>ping dest_IP source source_IP</code>	Verifies connectivity between source IP and destination IP
<code>router bgp AS-number</code>	Creates a BGP process and enters BGP process configuration mode
<code>router isis</code>	Creates an IS-IS process and enters IS-IS process configuration mode
<code>show ip bgp</code>	Displays BGP routing table
<code>show ip bgp summary</code>	Displays BGP routing protocol characteristics, including BGP neighbor status
<code>show ip interface brief</code>	Displays interface status and IPv4 addresses configured
<code>show ip route</code>	Displays the current routes in the routing table
<code>show isis database</code>	Displays the content of the IS-IS database
<code>show isis neighbors</code>	Displays IS-IS neighbor information
<code>show isis topology</code>	Displays lists of information related to the IS-IS topology for a specific router
<code>show running-config</code>	Displays running configuration

Cisco IOS XR Commands

Command	Description
<code>[no] shutdown</code>	Enables or disables interface on the router
<code>address-family ipv4 ipv6 unicast</code>	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)
<code>circuit-type level-1 level-1-2 level-2-only</code>	Enables IS-IS routing process to establish selected IS-IS circuit-type on the interface
<code>commit</code>	Commits changes to the running configuration
<code>configure terminal</code>	Enters configuration mode
<code>interface interface (global)</code>	Enters interface configuration mode
<code>interface interface (router)</code>	Defines the interfaces on which the IS-IS protocol runs
<code>ip address ip_address mask</code>	Sets the IPv4 address for an interface
<code>ipv6 enable</code>	Enables IPv6 support on the interface
<code>neighbor IP-address</code>	Configures BGP neighbor and enters BGP neighbor configuration mode
<code>network network/subnet-mask</code>	Enables BGP router to include network from routing table into BGP routing (BGP address family mode)
<code>next-hop-self</code>	Enables internal BGP router to send BGP routes with router BGP own IP address (BGP neighbor address family mode)
<code>pass</code>	Passes route for further processing (route-policy configuration mode)
<code>password clear password</code>	Enables authentication for BGP neighbor
<code>ping dest_IP source source_IP</code>	Verifies connectivity between source IP and destination IP (IPv4 and IPv6)
<code>remote-as AS-number</code>	Configures AS number for BGP neighbor. (BGP neighbor mode)
<code>route-policy route-policy-name</code>	Creates route policy and enters route policy configuration mode
<code>route-policy route-policy-name in out</code>	Applies route policy to BGP neighbor
<code>router bgp AS-number</code>	Creates a BGP process and enters BGP process configuration mode
<code>router isis process-ID</code>	Creates an IS-IS process
<code>show bgp</code>	Displays BGP routing table
<code>show bgp summary</code>	Displays BGP routing protocol characteristics, including BGP neighbor status
<code>show ipv4 interface brief</code>	Displays interface status and IPv4 addresses configured
<code>show isis database</code>	Displays the content of the IS-IS database
<code>show isis neighbors</code>	Displays IS-IS neighbor information

Command	Description
<code>show isis topology</code>	Displays lists of information related to the IS-IS topology for a specific router
<code>show route</code>	Displays the current routes in the routing table
<code>show running-config</code>	Displays running configuration
<code>update-source interface</code>	Enables BGP router to send BGP packets from source IP address (BGP neighbor configuration mode)

Task 1: Configure 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.

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 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 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. On the PE router, output should be similar to the following, taken from Pod 1:

```
RP/0/RSP0/CPU0:PE1#show bgp summary
< text omitted >
Neighbor      Spk    AS  MsgRcvd  MsgSent   TblVer  InQ  OutQ  Up/Down  St/PfxRcd
192.168.101.11  0 64501      17       14         1    0    0 00:11:56      0!
```

On the CE router, output should be similar to the following, taken from Pod 1:

```
CE1#show ip bgp summary
< text omitted >
Neighbor      V      AS  MsgRcvd  MsgSent   TblVer  InQ  OutQ  Up/Down  State/PfxRcd
192.168.101.10  4 64500      16       19         1    0    0 00:13:41      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.

Activity Verification

You have completed this task when you attain these results:

- 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”. On the PE router, output should be similar to the following, taken from Pod 1:

```
RP/0/RSP0/CPU0:PE1#show bgp summary
< text omitted >
Neighbor      Spk    AS  MsgRcvd  MsgSent   TblVer  InQ  OutQ  Up/Down  St/PfxRcd
192.168.101.11  0 64501      48       41         3    0    0 00:23:52      1
```

On the CE router, output should be similar to the following, taken from Pod 1:

```
CE1#show ip bgp summary
< text omitted >
Neighbor      V      AS  MsgRcvd  MsgSent   TblVer  InQ  OutQ  Up/Down  State/PfxRcd
192.168.101.10  4 64500      29       32         3    0    0 00:25:08      1
```

- 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. On the PE router, you should see the route from the CE router:

```
RP/0/RSP0/CPU0:PE1#show bgp
< text omitted >
      Network                Next Hop                Metric LocPrf Weight Path
*> 10.1.1.1/32                0.0.0.0                  0           32768 i
*> 10.1.10.1/32               192.168.101.11          0           0 64501 i
```

On the CE router, you should see the route from the PE router:

```
CE1#show ip bgp
< text omitted >
      Network                Next Hop                Metric LocPrf Weight Path
*> 10.1.1.1/32                192.168.101.10          0           0 64500 i
*> 10.1.10.1/32               0.0.0.0                  0           32768 i
```

Task 2: Configure 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 IPv4 and IPv6 address families.

On the PE router, you should see the interface toward the P1 router configured and running. Output should be similar to the following, taken from the PE1 router:

```
RP/0/RSP0/CPU0:PE1#show ipv4 interface brief | include Up
Loopback0                10.1.1.1                Up
MgmtEth0/RSP0/CPU0/0     10.10.10.17             Up
GigabitEthernet0/0/0/0    192.168.101.10          Up
GigabitEthernet0/0/0/1    192.168.112.10          Up
GigabitEthernet0/0/0/2    192.168.11.10           Up
GigabitEthernet0/0/0/3    192.168.12.10           Up
```

On the PE router, you should see the P1 router as an IS-IS neighbor. Output should be similar to the following, taken from the PE1 router:

```
RP/0/RSP0/CPU0:PE1#show isis neighbors
IS-IS 1 neighbors:
System Id      Interface      SNPA              State Holdtime Type IETF-NSF
CE1            Gi0/0/0/0      e8b7.482c.a180    Up    7        L1    Capable
P1             Gi0/0/0/2      4055.392e.d822    Up    7        L2    Capable
PE2            Gi0/0/0/1      e8b7.48fb.5801    Up    8        L2    Capable
```

Total neighbor count: 3

- 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.

Note The P1 router in the backbone is configured as the BGP route reflector.

On the PE router, you should see that an internal BGP session with P1 router starts. Output should be similar to the following, taken from the PE1 router:

```
RP/0/RSP0/CPU0:PE1#show bgp summary
< text omitted >
```

Neighbor St/PfxRcd	Spk	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down
10.0.1.1 6	0	64500	24	15	17	0	0	00:10:43
192.168.101.11 1	0	64501	31	42	17	0	0	00:15:51

On the PE router, you should see BGP routes coming from the P1 router. Verify that there are BGP routes from the neighbor pod. Output should be similar to the following, taken from the PE1 router:

```
RP/0/RSP0/CPU0:PE1#show bgp | include 10.2.1
*>i10.2.1.1/32          10.2.1.1          0      100      0 i
*>i10.2.10.1/32         192.168.102.21      0      100      0 64502 i
```

As you can see, the next-hop IP address is different for both routes coming from the neighbor pod.

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.

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

Activity Verification

You have completed this task when you attain these results:

- On the PE router in your pod, you should see BGP routes from the neighbor pod, with next-hop IP addresses set to the neighbor PE router Loopback0 interface IP address. Output should be similar to the following, taken from the PE1 router, where 10.2.1.1 is the IP address from the PE2 Loopback0 interface:

```
RP/0/RSP0/CPU0:PE1#show bgp | in 10.2.1
*>i10.2.1.1/32          10.2.1.1          0      100      0 i
*>i10.2.10.1/32         10.2.1.1          0      100      0 64502 i
```

The neighbor pod PE router should have routes coming from your pod PE router with your PE router Loopback0 interface IP address. Output should be similar to the following, taken from the PE2 router, where 10.1.1.1 is the IP address from the PE1 Loopback0 interface:

```
PE2#show ip bgp | include 10.1.1
r>i10.1.1.1/32          10.1.1.1          0      100      0 i
*>i10.1.10.1/32         10.1.1.1          0      100      0 64501 i
```

- 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.

```
CE1#ping 10.2.10.1 (IPv4 address of the CE2 router Loopback0 interface)
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.2.10.1, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

Lab 5-1: Implement Route Redistribution

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

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 sent to the pod CE router by configuring one-way redistribution.

Note	Students from two different pods are working in 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.

Visual Objective for Lab 5-1: Implement Route Redistribution



© 2012 Cisco and/or its affiliates. All rights reserved.

SPROUTE v1.01-12

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

Cisco IOS XR Commands

Command	Description
<code>address-family ipv4 ipv6 unicast</code>	Enables IPv4 or IPv6 IS-IS or OSPF routing and enters address family configuration mode
<code>area area</code>	(Router) Creates OSPF area on the router
<code>commit</code>	Commits changes to the running configuration
<code>configure terminal</code>	Enters configuration mode
<code>default-information originate always</code>	Originates OSPF default route, even if there is no default route in the routing table (always)
<code>distance ospf external value</code>	Modifies OSPF external administrative distance
<code>if condition then action-1 else action-2 endif</code>	Routes policy list if sentence
<code>interface interface</code>	(Router) Defines the interfaces on which the OSPF or IS-IS protocol runs
<code>pass</code>	Passes route for further processing (route-policy configuration mode)
<code>ping dest_IP source source_IP</code>	Verifies connectivity between source IP and destination IP (IPv4 and IPv6)
<code>redistribute protocol</code>	Redistributes protocol into OSPF or IS-IS
<code>route-policy route-policy-name</code>	Creates route policy and enters route policy configuration mode
<code>router isis process-ID</code>	Creates an IS-IS process and enters router IS-IS process configuration mode
<code>router ospf ospfv3 process-ID</code>	Creates an OSPFv2 or OSPFv3 process and enters router OSPFv2 or OSPFv3 process configuration mode
<code>show ipv4 interface brief</code>	Displays interface status and IPv4 addresses configured
<code>show isis database</code>	Displays the content of the IS-IS database
<code>show isis neighbors</code>	Displays IS-IS neighbor information
<code>show isis topology</code>	Displays lists of information related to the IS-IS topology for a specific router
<code>show ospf</code>	Displays information related to the OSPF routing protocol running on the router
<code>show ospf ospfv3 database</code>	Displays the content of the OSPF database
<code>show ospf ospfv3 interface</code>	Displays OSPF interface information
<code>show ospf ospfv3 neighbors</code>	Displays OSPF neighbor information
<code>show protocols</code>	Displays protocols running on the router
<code>show route</code>	Displays the current routes in the routing table
<code>show running-config</code>	Displays running configuration

Task 1: Redistribute 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.

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 an 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 an OSPF process 1. Start an 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. Check that from the pod CE router, you are not able to ping core devices.

On the PE router, **show** command output should be similar to the following, taken from PE1 (Cisco IOS XR) router:

```
RP/0/RSP0/CPU0:PE1#show isis neighbors
```

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

System Id	Interface	SNPA	State	Holdtime	Type	IETF-NSF
P1	Gi0/0/0/2	4055.392e.d822	Up	9	L2	Capable

```
Total neighbor count: 1
RP/0/RSP0/CPU0:PE1#show ospf neighbor
* Indicates MADJ interface
```

Neighbors for OSPF 1

Neighbor ID	Pri	State	Dead Time	Address	Interface
10.2.1.1	1	FULL/BDR	00:00:37	192.168.112.20	GigabitEthernet0/0/0/1
Neighbor is up for 01:14:35					
10.1.10.1	1	FULL/DR	00:00:33	192.168.101.11	GigabitEthernet0/0/0/0
Neighbor is up for 01:18:31					

Total neighbor count: 2

Following outputs are taken from PE2 (Cisco IOS XE) router:

```
PE2#show isis neighbors
```

```
Tag null:
```

System Id	Type	Interface	IP Address	State	Holdtime	Circuit Id
P1	L2	Gi0/0/2	192.168.21.1	UP	24	PE2.04

```
PE2#show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
10.1.1.1	1	FULL/DR	00:00:37	192.168.112.10	GigabitEthernet0/0/1
10.2.10.1	1	FULL/DR	00:00:38	192.168.102.21	GigabitEthernet0/0/0

On the CE router, **show** command output should be similar to the following, taken from pods 1 and 2. There should be a summary LSA for Loopback0 from the neighbor pod CE router present in the OSPF database:

```
CE1#show ip ospf database
```

OSPF Router with ID (10.1.10.1) (Process ID 1)

Router Link States (Area 11)

Link ID	ADV Router	Age	Seq#	Checksum	Link count
10.1.1.1	10.1.1.1	926	0x80000004	0x004220	1
10.1.10.1	10.1.10.1	928	0x80000005	0x0049DE	2

Net Link States (Area 11)

Link ID	ADV Router	Age	Seq#	Checksum
192.168.101.11	10.1.10.1	928	0x80000003	0x00B96B

Summary Net Link States (Area 11)

Link ID	ADV Router	Age	Seq#	Checksum
10.2.10.1	10.1.1.1	671	0x80000003	0x0021F2
192.168.102.0	10.1.1.1	672	0x80000003	0x00124A
192.168.112.0	10.1.1.1	927	0x80000003	0x0099B9

!

```
CE2#show ip ospf database
```

OSPF Router with ID (10.2.10.1) (Process ID 1)

Router Link States (Area 12)

Link ID	ADV Router	Age	Seq#	Checksum	Link count
10.2.1.1	10.2.1.1	751	0x80000004	0x003515	1
10.2.10.1	10.2.10.1	835	0x80000005	0x004DC1	2

Net Link States (Area 12)

Link ID	ADV Router	Age	Seq#	Checksum
192.168.102.21	10.2.10.1	835	0x80000003	0x0056C0

Summary Net Link States (Area 12)

Link ID	ADV Router	Age	Seq#	Checksum
10.1.10.1	10.2.1.1	751	0x80000003	0x0025EE
192.168.101.0	10.2.1.1	751	0x80000003	0x001547
192.168.112.0	10.2.1.1	751	0x80000003	0x0091C0

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

```
CE1#ping 10.0.1.1 (P1 Loopback0 IP address)
```

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 IS-IS and OSPF routing protocols. Redistribute only IS-IS Level 2 routes into OSPF and make OSPF routes redistributed into IS-IS as Level 2.

Activity Verification

You have completed this task when you attain this result:

- On the pod PE router, on the OSPF database, you should see routes redistributed from IS-IS routing protocol as external routes. The following outputs are taken from pod 1 and pod 2:

```
RP/0/RSP0/CPU0:PE1#show ospf database
```

< text omitted >

Type-5 AS External Link States

Link ID	ADV Router	Age	Seq#	Checksum	Tag
10.0.1.1	10.1.1.1	896	0x80000001	0x00e0a6	0
10.0.2.1	10.1.1.1	896	0x80000001	0x00d5b0	0
10.1.1.1	10.1.1.1	896	0x80000001	0x00d4b1	0
10.2.1.1	10.1.1.1	896	0x80000001	0x00c8bc	0
192.168.1.0	10.1.1.1	896	0x80000001	0x00bb6d	0
192.168.2.0	10.1.1.1	896	0x80000001	0x00b077	0
192.168.11.0	10.1.1.1	896	0x80000001	0x004dd1	0
192.168.12.0	10.2.1.1	714	0x80000001	0x003ae2	0
192.168.21.0	10.1.1.1	896	0x80000001	0x00de36	0
192.168.22.0	10.1.1.1	896	0x80000001	0x00d340	0
192.168.31.0	10.1.1.1	896	0x80000001	0x00709a	0
192.168.42.0	10.1.1.1	896	0x80000001	0x00f609	0
192.168.51.0	10.1.1.1	896	0x80000001	0x009363	0
192.168.52.0	10.1.1.1	896	0x80000001	0x00886d	0
192.168.61.0	10.1.1.1	896	0x80000001	0x0025c7	0
192.168.62.0	10.1.1.1	896	0x80000001	0x001ad1	0
192.168.71.0	10.1.1.1	896	0x80000001	0x00b62c	0
192.168.82.0	10.1.1.1	896	0x80000001	0x003d9a	0

!

```
PE2#show ip ospf database
```

< text omitted >

Type-5 AS External Link States

Link ID	ADV Router	Age	Seq#	Checksum	Tag
10.0.1.1	10.1.1.1	985	0x80000001	0x00E0A6	0
10.0.2.1	10.1.1.1	985	0x80000001	0x00D5B0	0
10.1.1.1	10.1.1.1	985	0x80000001	0x00D4B1	0
10.2.1.1	10.1.1.1	985	0x80000001	0x00C8BC	0
192.168.1.0	10.1.1.1	985	0x80000001	0x00BB6D	0
192.168.2.0	10.1.1.1	985	0x80000001	0x00B077	0
192.168.11.0	10.1.1.1	985	0x80000001	0x004DD1	0

192.168.12.0	10.2.1.1	800	0x80000001	0x003AE2	0
192.168.21.0	10.1.1.1	985	0x80000001	0x00DE36	0
192.168.22.0	10.1.1.1	985	0x80000001	0x00D340	0
192.168.31.0	10.1.1.1	985	0x80000001	0x00709A	0
192.168.42.0	10.1.1.1	985	0x80000001	0x00F609	0
192.168.51.0	10.1.1.1	985	0x80000001	0x009363	0
192.168.52.0	10.1.1.1	985	0x80000001	0x00886D	0
192.168.61.0	10.1.1.1	985	0x80000001	0x0025C7	0
192.168.62.0	10.1.1.1	985	0x80000001	0x001AD1	0
192.168.71.0	10.1.1.1	985	0x80000001	0x00B62C	0
192.168.82.0	10.1.1.1	985	0x80000001	0x003D9A	0

Task 2: Prevent 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 OSPFintoISIS.
- Step 2** On the pod PE router, apply the RPL or route map configured in the previous step to the OSPF into IS-IS redistribution configuration.
- Step 3** On the pod PE router, configure 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 configured in the previous step to the IS-IS into OSPF redistribution configuration.

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

Visual Objective for Lab 5-1 Task 2: Prevent Potential Routing Loops



© 2012 Cisco and/or its affiliates. All rights reserved.

SPROUTE v1.01-14

Activity Verification

You have completed this task when you attain these results:

- On the pod PE router, on the OSPF database, you should see routes redistributed from IS-IS routing protocol as external routes with tag 115. The following outputs are taken from pods 1 and 2:

RP/0/RSP0/CPU0:PE1#show ospf database

< text omitted >

Type-5 AS External Link States

Link ID	ADV Router	Age	Seq#	Checksum	Tag
10.0.1.1	10.1.1.1	631	0x80000003	0x00fa17	115
10.0.2.1	10.1.1.1	631	0x80000003	0x00ef21	115
10.1.1.1	10.1.1.1	631	0x80000004	0x00ec23	115
10.2.1.1	10.1.1.1	631	0x80000004	0x00e02e	115
192.168.1.0	10.1.1.1	631	0x80000003	0x00d5dd	115
192.168.2.0	10.1.1.1	631	0x80000003	0x00cae7	115
192.168.11.0	10.1.1.1	631	0x80000004	0x006543	115
192.168.12.0	10.2.1.1	982	0x80000002	0x005652	115
192.168.21.0	10.1.1.1	631	0x80000004	0x00f6a7	115
192.168.22.0	10.1.1.1	631	0x80000004	0x00ebb1	115
192.168.31.0	10.1.1.1	631	0x80000003	0x008a0b	115
192.168.42.0	10.1.1.1	631	0x80000003	0x001179	115
192.168.51.0	10.1.1.1	631	0x80000003	0x00add3	115
192.168.52.0	10.1.1.1	631	0x80000003	0x00a2dd	115
192.168.61.0	10.1.1.1	631	0x80000003	0x003f38	115
192.168.62.0	10.1.1.1	631	0x80000003	0x003442	115
192.168.71.0	10.1.1.1	631	0x80000003	0x00d09c	115
192.168.82.0	10.1.1.1	631	0x80000003	0x00570b	115

!

PE2#show ip ospf database

< text omitted >

Type-5 AS External Link States

Link ID	ADV Router	Age	Seq#	Checksum	Tag
10.0.1.1	10.1.1.1	677	0x80000003	0x00FA17	115
10.0.2.1	10.1.1.1	677	0x80000003	0x00EF21	115
10.1.1.1	10.1.1.1	677	0x80000004	0x00EC23	115

10.2.1.1	10.1.1.1	677	0x80000004	0x00E02E	115
192.168.1.0	10.1.1.1	677	0x80000003	0x00D5DD	115
192.168.2.0	10.1.1.1	677	0x80000003	0x00CAE7	115
192.168.11.0	10.1.1.1	677	0x80000004	0x006543	115
192.168.12.0	10.2.1.1	1026	0x80000002	0x005652	115
192.168.21.0	10.1.1.1	677	0x80000004	0x00F6A7	115
192.168.22.0	10.1.1.1	677	0x80000004	0x00EBB1	115
192.168.31.0	10.1.1.1	677	0x80000003	0x008A0B	115
192.168.42.0	10.1.1.1	677	0x80000003	0x001179	115
192.168.51.0	10.1.1.1	677	0x80000003	0x00ADD3	115
192.168.52.0	10.1.1.1	677	0x80000003	0x00A2DD	115
192.168.61.0	10.1.1.1	677	0x80000003	0x003F38	115
192.168.62.0	10.1.1.1	677	0x80000003	0x003442	115
192.168.71.0	10.1.1.1	677	0x80000003	0x00D09C	115
192.168.82.0	10.1.1.1	677	0x80000003	0x00570B	115

- 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 same routes as OSPF external routes. This is not optimal routing, since one PE router is using PE to PE link to reach networks in the backbone. The following outputs are taken from pods 1 and 2:

```
RP/0/RSP0/CPU0:PE1#show route
```

```
< text omitted >
```

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

```
i L2 10.0.1.1/32 [115/20] via 192.168.11.1, 00:52:09, GigabitEthernet0/0/0/2
i L2 10.0.2.1/32 [115/30] via 192.168.11.1, 00:52:09, GigabitEthernet0/0/0/2
L 10.1.1.1/32 is directly connected, 03:30:14, Loopback0
B 10.1.10.1/32 [20/0] via 192.168.101.11, 03:28:03
i L2 10.2.1.1/32 [115/30] via 192.168.11.1, 03:10:41, GigabitEthernet0/0/0/2
O IA 10.2.10.1/32 [110/3] via 192.168.112.20, 00:57:51, GigabitEthernet0/0/0/1
C 10.10.10.0/24 is directly connected, 03:30:14, MgmtEth0/RSP0/CPU0/0
L 10.10.10.17/32 is directly connected, 03:30:14, MgmtEth0/RSP0/CPU0/0
i L2 192.168.1.0/24 [115/20] via 192.168.11.1, 00:52:09,
GigabitEthernet0/0/0/2
i L2 192.168.2.0/24 [115/20] via 192.168.11.1, 00:52:09,
GigabitEthernet0/0/0/2
C 192.168.11.0/24 is directly connected, 03:30:12, GigabitEthernet0/0/0/2
L 192.168.11.10/32 is directly connected, 03:30:12, GigabitEthernet0/0/0/2
C 192.168.12.0/24 is directly connected, 03:30:12, GigabitEthernet0/0/0/3
L 192.168.12.10/32 is directly connected, 03:30:12, GigabitEthernet0/0/0/3
i L2 192.168.21.0/24 [115/20] via 192.168.11.1, 03:10:41,
GigabitEthernet0/0/0/2
i L2 192.168.22.0/24 [115/30] via 192.168.11.1, 03:30:05,
GigabitEthernet0/0/0/2
i L2 192.168.31.0/24 [115/20] via 192.168.11.1, 00:52:09,
GigabitEthernet0/0/0/2
i L2 192.168.42.0/24 [115/30] via 192.168.11.1, 00:52:09,
GigabitEthernet0/0/0/2
i L2 192.168.51.0/24 [115/20] via 192.168.11.1, 00:52:09,
GigabitEthernet0/0/0/2
i L2 192.168.52.0/24 [115/30] via 192.168.11.1, 00:52:09,
GigabitEthernet0/0/0/2
i L2 192.168.61.0/24 [115/20] via 192.168.11.1, 00:52:09,
GigabitEthernet0/0/0/2
i L2 192.168.62.0/24 [115/30] via 192.168.11.1, 00:52:09,
GigabitEthernet0/0/0/2
i L2 192.168.71.0/24 [115/20] via 192.168.11.1, 00:52:09,
GigabitEthernet0/0/0/2
i L2 192.168.82.0/24 [115/30] via 192.168.11.1, 00:52:09,
GigabitEthernet0/0/0/2
C 192.168.101.0/24 is directly connected, 03:30:09, GigabitEthernet0/0/0/0
L 192.168.101.10/32 is directly connected, 03:30:09, GigabitEthernet0/0/0/0
O IA 192.168.102.0/24 [110/2] via 192.168.112.20, 00:57:51,
GigabitEthernet0/0/0/1
C 192.168.112.0/24 is directly connected, 00:58:33, GigabitEthernet0/0/0/1
L 192.168.112.10/32 is directly connected, 00:58:33, GigabitEthernet0/0/0/1
!
```

```

PE2#show ip route
< text 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 192.168.112.10, 00:52:31, GigabitEthernet0/0/1
O E2   10.0.2.1 [110/20] via 192.168.112.10, 00:52:31, GigabitEthernet0/0/1
O E2   10.1.1.1 [110/20] via 192.168.112.10, 00:52:31, GigabitEthernet0/0/1
O IA   10.1.10.1 [110/3] via 192.168.112.10, 00:58:07, GigabitEthernet0/0/1
C      10.2.1.1 is directly connected, Loopback0
B      10.2.10.1 [20/0] via 192.168.102.21, 03:30:35
O E2   192.168.1.0/24
        [110/20] via 192.168.112.10, 00:52:31, GigabitEthernet0/0/1
O E2   192.168.2.0/24
        [110/20] via 192.168.112.10, 00:52:31, GigabitEthernet0/0/1
O E2   192.168.11.0/24
        [110/20] via 192.168.112.10, 00:52:31, GigabitEthernet0/0/1
i L2   192.168.12.0/24
        [115/30] via 192.168.21.1, 03:31:39, GigabitEthernet0/0/2
        192.168.21.0/24 is variably subnetted, 2 subnets, 2 masks
C      192.168.21.0/24 is directly connected, GigabitEthernet0/0/2
L      192.168.21.20/32 is directly connected, GigabitEthernet0/0/2
        192.168.22.0/24 is variably subnetted, 2 subnets, 2 masks
C      192.168.22.0/24 is directly connected, GigabitEthernet0/0/3
L      192.168.22.20/32 is directly connected, GigabitEthernet0/0/3
O E2   192.168.31.0/24
        [110/20] via 192.168.112.10, 00:52:31, GigabitEthernet0/0/1
O E2   192.168.42.0/24
        [110/20] via 192.168.112.10, 00:52:31, GigabitEthernet0/0/1
O E2   192.168.51.0/24
        [110/20] via 192.168.112.10, 00:52:31, GigabitEthernet0/0/1
O E2   192.168.52.0/24
        [110/20] via 192.168.112.10, 00:52:31, GigabitEthernet0/0/1
O E2   192.168.61.0/24
        [110/20] via 192.168.112.10, 00:52:31, GigabitEthernet0/0/1
O E2   192.168.62.0/24
        [110/20] via 192.168.112.10, 00:52:31, GigabitEthernet0/0/1
O E2   192.168.71.0/24
        [110/20] via 192.168.112.10, 00:52:31, GigabitEthernet0/0/1
O E2   192.168.82.0/24
        [110/20] via 192.168.112.10, 00:52:31, GigabitEthernet0/0/1
O IA   192.168.101.0/24
        [110/2] via 192.168.112.10, 00:58:07, GigabitEthernet0/0/1
        192.168.102.0/24 is variably subnetted, 2 subnets, 2 masks
C      192.168.102.0/24 is directly connected, GigabitEthernet0/0/0
L      192.168.102.20/32 is directly connected, GigabitEthernet0/0/0
        192.168.112.0/24 is variably subnetted, 2 subnets, 2 masks
C      192.168.112.0/24 is directly connected, GigabitEthernet0/0/1
L      192.168.112.20/32 is directly connected, GigabitEthernet0/0/1

```

Task 3: Modify 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 these steps:

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

Activity Verification

Complete lab activity verification:

- On the pod PE router verify the routing table. Both PE routers in the team will show IS-IS routes for backbone networks, representing optimal routing. The following outputs are taken from pod 1 and pod 2:

```
RP/0/RSP0/CPU0:PE1#show route
```

```
< text omitted >
```

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

```
i L2 10.0.1.1/32 [115/20] via 192.168.11.1, 01:24:54, GigabitEthernet0/0/0/2
i L2 10.0.2.1/32 [115/30] via 192.168.11.1, 01:24:54, GigabitEthernet0/0/0/2
L 10.1.1.1/32 is directly connected, 04:02:58, Loopback0
B 10.1.10.1/32 [20/0] via 192.168.101.11, 04:00:47
i L2 10.2.1.1/32 [115/30] via 192.168.11.1, 03:43:25, GigabitEthernet0/0/0/2
O IA 10.2.10.1/32 [110/3] via 192.168.112.20, 00:06:36, GigabitEthernet0/0/0/1
C 10.10.10.0/24 is directly connected, 04:02:58, MgmtEth0/RSP0/CPU0/0
L 10.10.10.17/32 is directly connected, 04:02:58, MgmtEth0/RSP0/CPU0/0
i L2 192.168.1.0/24 [115/20] via 192.168.11.1, 01:24:54,
GigabitEthernet0/0/0/2
i L2 192.168.2.0/24 [115/20] via 192.168.11.1, 01:24:54,
GigabitEthernet0/0/0/2
C 192.168.11.0/24 is directly connected, 04:02:57, GigabitEthernet0/0/0/2
L 192.168.11.10/32 is directly connected, 04:02:57, GigabitEthernet0/0/0/2
C 192.168.12.0/24 is directly connected, 04:02:57, GigabitEthernet0/0/0/3
L 192.168.12.10/32 is directly connected, 04:02:57, GigabitEthernet0/0/0/3
i L2 192.168.21.0/24 [115/20] via 192.168.11.1, 03:43:25,
GigabitEthernet0/0/0/2
i L2 192.168.22.0/24 [115/30] via 192.168.11.1, 04:02:49,
GigabitEthernet0/0/0/2
i L2 192.168.31.0/24 [115/20] via 192.168.11.1, 01:24:54,
GigabitEthernet0/0/0/2
i L2 192.168.42.0/24 [115/30] via 192.168.11.1, 01:24:54,
GigabitEthernet0/0/0/2
i L2 192.168.51.0/24 [115/20] via 192.168.11.1, 01:24:54,
GigabitEthernet0/0/0/2
i L2 192.168.52.0/24 [115/30] via 192.168.11.1, 01:24:54,
GigabitEthernet0/0/0/2
i L2 192.168.61.0/24 [115/20] via 192.168.11.1, 01:24:54,
GigabitEthernet0/0/0/2
i L2 192.168.62.0/24 [115/30] via 192.168.11.1, 01:24:54,
GigabitEthernet0/0/0/2
i L2 192.168.71.0/24 [115/20] via 192.168.11.1, 01:24:54,
GigabitEthernet0/0/0/2
i L2 192.168.82.0/24 [115/30] via 192.168.11.1, 01:24:54,
GigabitEthernet0/0/0/2
C 192.168.101.0/24 is directly connected, 04:02:54, GigabitEthernet0/0/0/0
L 192.168.101.10/32 is directly connected, 04:02:54, GigabitEthernet0/0/0/0
O IA 192.168.102.0/24 [110/2] via 192.168.112.20, 00:06:36,
GigabitEthernet0/0/0/1
C 192.168.112.0/24 is directly connected, 01:31:18, GigabitEthernet0/0/0/1
L 192.168.112.10/32 is directly connected, 01:31:18, GigabitEthernet0/0/0/1
!
```

```
PE2#show ip route
```

```
< text 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 192.168.21.1, 00:06:24, GigabitEthernet0/0/2
i L2 10.0.2.1 [115/30] via 192.168.21.1, 00:06:24, GigabitEthernet0/0/2
i L2 10.1.1.1 [115/30] via 192.168.21.1, 00:06:24, GigabitEthernet0/0/2
O IA 10.1.10.1 [110/3] via 192.168.112.10, 00:06:24, GigabitEthernet0/0/1
C 10.2.1.1 is directly connected, Loopback0
B 10.2.10.1 [20/0] via 192.168.102.21, 04:03:12
i L2 192.168.1.0/24 [115/20] via 192.168.21.1, 00:06:24, GigabitEthernet0/0/2
i L2 192.168.2.0/24 [115/20] via 192.168.21.1, 00:06:24, GigabitEthernet0/0/2
i L2 192.168.11.0/24
[115/20] via 192.168.21.1, 00:06:24, GigabitEthernet0/0/2
i L2 192.168.12.0/24
[115/30] via 192.168.21.1, 04:04:16, GigabitEthernet0/0/2
```

```

192.168.21.0/24 is variably subnetted, 2 subnets, 2 masks
C    192.168.21.0/24 is directly connected, GigabitEthernet0/0/2
L    192.168.21.20/32 is directly connected, GigabitEthernet0/0/2
192.168.22.0/24 is variably subnetted, 2 subnets, 2 masks
C    192.168.22.0/24 is directly connected, GigabitEthernet0/0/3
L    192.168.22.20/32 is directly connected, GigabitEthernet0/0/3
i L2 192.168.31.0/24
      [115/20] via 192.168.21.1, 00:06:24, GigabitEthernet0/0/2
i L2 192.168.42.0/24
      [115/30] via 192.168.21.1, 00:06:24, GigabitEthernet0/0/2
i L2 192.168.51.0/24
      [115/20] via 192.168.21.1, 00:06:24, GigabitEthernet0/0/2
i L2 192.168.52.0/24
      [115/30] via 192.168.21.1, 00:06:24, GigabitEthernet0/0/2
i L2 192.168.61.0/24
      [115/20] via 192.168.21.1, 00:06:24, GigabitEthernet0/0/2
i L2 192.168.62.0/24
      [115/30] via 192.168.21.1, 00:06:24, GigabitEthernet0/0/2
i L2 192.168.71.0/24
      [115/20] via 192.168.21.1, 00:06:24, GigabitEthernet0/0/2
i L2 192.168.82.0/24
      [115/30] via 192.168.21.1, 00:06:24, GigabitEthernet0/0/2
O IA 192.168.101.0/24
      [110/2] via 192.168.112.10, 00:06:24, GigabitEthernet0/0/1
192.168.102.0/24 is variably subnetted, 2 subnets, 2 masks
C    192.168.102.0/24 is directly connected, GigabitEthernet0/0/0
L    192.168.102.20/32 is directly connected, GigabitEthernet0/0/0
192.168.112.0/24 is variably subnetted, 2 subnets, 2 masks
C    192.168.112.0/24 is directly connected, GigabitEthernet0/0/1
L    192.168.112.20/32 is directly connected, GigabitEthernet0/0/1

```

- On the pod CE router verify the routing table, you should see external OSPF routes for all backbone networks. The following outputs are taken from pods 1 and 2:

```
CE1#show ip route ospf
```

```
< text 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 192.168.101.10, 01:30:10, GigabitEthernet0/0
O E2 10.0.2.1 [110/20] via 192.168.101.10, 01:30:10, GigabitEthernet0/0
O E2 192.168.1.0/24 [110/20] via 192.168.101.10, 01:30:10, GigabitEthernet0/0
O E2 192.168.2.0/24 [110/20] via 192.168.101.10, 01:30:10, GigabitEthernet0/0
O E2 192.168.11.0/24
      [110/20] via 192.168.101.10, 01:30:10, GigabitEthernet0/0
O E2 192.168.12.0/24
      [110/20] via 192.168.101.10, 01:33:43, GigabitEthernet0/0
O E2 192.168.21.0/24
      [110/20] via 192.168.101.10, 01:30:11, GigabitEthernet0/0
O E2 192.168.22.0/24
      [110/20] via 192.168.101.10, 01:30:11, GigabitEthernet0/0
O E2 192.168.31.0/24
      [110/20] via 192.168.101.10, 01:30:12, GigabitEthernet0/0
O E2 192.168.42.0/24
      [110/20] via 192.168.101.10, 01:30:12, GigabitEthernet0/0
O E2 192.168.51.0/24
      [110/20] via 192.168.101.10, 01:30:12, GigabitEthernet0/0
O E2 192.168.52.0/24
      [110/20] via 192.168.101.10, 01:30:12, GigabitEthernet0/0
O E2 192.168.61.0/24
      [110/20] via 192.168.101.10, 01:30:12, GigabitEthernet0/0
O E2 192.168.62.0/24
      [110/20] via 192.168.101.10, 01:30:12, GigabitEthernet0/0
O E2 192.168.71.0/24
      [110/20] via 192.168.101.10, 01:30:12, GigabitEthernet0/0
O E2 192.168.82.0/24
      [110/20] via 192.168.101.10, 01:30:12, GigabitEthernet0/0
O IA 192.168.102.0/24

```

```

[110/3] via 192.168.101.10, 01:35:53, GigabitEthernet0/0
O IA 192.168.112.0/24
[110/2] via 192.168.101.10, 01:35:48, GigabitEthernet0/0
!
CE2#show ip route ospf
< text 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 192.168.102.20, 00:18:17, GigabitEthernet0/0
O E2 10.0.2.1 [110/20] via 192.168.102.20, 00:18:17, GigabitEthernet0/0
O E2 192.168.1.0/24 [110/20] via 192.168.102.20, 00:18:17, GigabitEthernet0/0
O E2 192.168.2.0/24 [110/20] via 192.168.102.20, 00:18:17, GigabitEthernet0/0
O E2 192.168.11.0/24
[110/20] via 192.168.102.20, 00:18:17, GigabitEthernet0/0
O E2 192.168.12.0/24
[110/20] via 192.168.102.20, 01:40:34, GigabitEthernet0/0
O E2 192.168.21.0/24
[110/20] via 192.168.102.20, 01:37:01, GigabitEthernet0/0
O E2 192.168.22.0/24
[110/20] via 192.168.102.20, 01:37:01, GigabitEthernet0/0
O E2 192.168.31.0/24
[110/20] via 192.168.102.20, 00:18:18, GigabitEthernet0/0
O E2 192.168.42.0/24
[110/20] via 192.168.102.20, 00:18:18, GigabitEthernet0/0
O E2 192.168.51.0/24
[110/20] via 192.168.102.20, 00:18:18, GigabitEthernet0/0
O E2 192.168.52.0/24
[110/20] via 192.168.102.20, 00:18:18, GigabitEthernet0/0
O E2 192.168.61.0/24
[110/20] via 192.168.102.20, 00:18:18, GigabitEthernet0/0
O E2 192.168.62.0/24
[110/20] via 192.168.102.20, 00:18:19, GigabitEthernet0/0
O E2 192.168.71.0/24
[110/20] via 192.168.102.20, 00:18:19, GigabitEthernet0/0
O E2 192.168.82.0/24
[110/20] via 192.168.102.20, 00:18:19, GigabitEthernet0/0
O IA 192.168.101.0/24
[110/3] via 192.168.102.20, 01:42:40, GigabitEthernet0/0
O IA 192.168.112.0/24
[110/2] via 192.168.102.20, 01:43:15, 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 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.

Activity Verification

Complete lab activity verification:

- On the pod CE router, on the routing table, you should see the default OSPF route pointing to the pod PE router. The following outputs are taken from pods 1 and 2:

```

CE1#show ip route ospf
< text omitted >
Gateway of last resort is 192.168.101.10 to network 0.0.0.0

O*E2  0.0.0.0/0 [110/1] via 192.168.101.10, 00:04:37, GigabitEthernet0/0
O IA   192.168.102.0/24
        [110/3] via 192.168.101.10, 01:58:05, GigabitEthernet0/0
O IA   192.168.112.0/24
        [110/2] via 192.168.101.10, 01:58:00, GigabitEthernet0/0
!
CE2#show ip route ospf
< text omitted >
Gateway of last resort is 192.168.102.20 to network 0.0.0.0

O*E2  0.0.0.0/0 [110/1] via 192.168.102.20, 00:04:20, GigabitEthernet0/0
O IA   192.168.101.0/24
        [110/3] via 192.168.102.20, 01:58:25, GigabitEthernet0/0
O IA   192.168.112.0/24
        [110/2] via 192.168.102.20, 01:59:00, GigabitEthernet0/0

```

- 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. The following output is taken from pod 1:

```

CE1#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 = 1/1/4 ms

```

Lab 5-2: Influence BGP Route Selection

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

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.


Note Students from two different pods are working in team. The CE routers in both pods are running Cisco IOS Software. First pod in the team will work on the PE router running Cisco IOS XR Software and 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 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
<code>[no] ip ipv6 router isis</code>	Enables or disables IS-IS routing to the interface
<code>[no] isis circuit-type level-1 level-1-2 level-2-only</code>	Enables or disables IS-IS routing process to establish selected IS-IS circuit type on the interface
<code>router isis</code>	Creates or enters an IS-IS process
<code>[no] shutdown</code>	Enables or disables interface on the router
<code>ip bgp-community new-format</code>	Enables BGP community new format presentation
<code>ip community-list expanded name permit deny value</code>	Creates community list
<code>ip ipv6 router isis</code>	Enables IS-IS routing to the interface
<code>isis circuit-type level-1 level-1-2 level-2-only</code>	Enables IS-IS routing process to establish selected IS-IS circuit type on the interface
<code>neighbor IP-address ebgp-multihop</code>	Configures EBGp multihop feature for BGP neighbor
<code>neighbor IP-address remote-as AS-number</code>	Configures BGP neighbor
<code>neighbor IP-address route-map name in out</code>	Applies route map to the BGP neighbor
<code>neighbor IP-address send-community both</code>	Enables router to send standard and extended BGP communities to the BGP neighbor
<code>neighbor IP-address weight value</code>	Configures BGP weight for routes received from BGP neighbor
<code>route-map name</code>	Creates and enters route map
<code>router bgp AS-number</code>	Creates a BGP process and enters BGP process configuration mode
<code>router isis</code>	Creates an IS-IS process and enters IS-IS process configuration mode
<code>set match tag local-preference metric community value</code>	Route map set or match option
<code>show ip bgp</code>	Displays BGP routing table
<code>show ip bgp summary</code>	Displays BGP routing protocol characteristics, including BGP neighbor status
<code>show isis database neighbors topology</code>	Displays the content of the IS-IS database or neighbor information or displays lists of information related to the IS-IS topology for a specific router


Cisco IOS XR Commands

Command	Description
<code>[no] shutdown</code>	Enables or disables interface on the router
<code>address-family ipv4 ipv6 unicast</code>	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)
<code>circuit-type level-1 level-1-2 level-2-only</code>	Enables IS-IS routing process to establish selected IS-IS circuit type on the interface
<code>commit</code>	Commits changes to the running configuration
<code>ebgp-multihop</code>	Configures EBGp multihop feature for BGP neighbor
<code>if condition then action-1 else action-2 endif</code>	Routes policy list if sentence
<code>interface interface (router)</code>	Defines the interfaces on which the IS-IS protocol runs
<code>neighbor IP-address</code>	Configures BGP neighbor and enters BGP neighbor configuration mode
<code>pass</code>	Passes route for further processing (route-policy configuration mode)
<code>prepend as-path value repets</code>	Enables AS-path prepending in the RPL
<code>remote-as AS-number</code>	Configures AS number for BGP neighbor (BGP neighbor mode)
<code>route-policy route-policy-name</code>	Creates route policy and enters route policy configuration mode
<code>route-policy route-policy-name in out</code>	Applies route policy to BGP neighbor
<code>router bgp AS-number</code>	Creates a BGP process and enters BGP process configuration mode
<code>router isis process-ID</code>	Creates an IS-IS process and enters router IS-IS process configuration mode
<code>set local-preference med value</code>	Sets BGP local preference or MED in the RPL
<code>show bgp</code>	Displays BGP routing table
<code>show bgp summary</code>	Displays BGP routing protocol characteristics, including BGP neighbor status
<code>show isis database neighbors topology</code>	Displays the content of the IS-IS database or neighbor information or displays lists of information related to the IS-IS topology for a specific router
<code>show route</code>	Displays the current routes in the routing table

Task 0: Set Second Link Between CE and PE

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

The figure shows what you will accomplish in this task:



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.

Activity Verification

You have completed this task when you attain this result:

- On the pod CE router, test connectivity to the neighbor pod PE router. The following outputs are taken from the pod CE routers:

```
CE1#ping 192.168.201.20
```

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

```
Sending 5, 100-byte ICMP Echos to 192.168.201.20, timeout is 2 seconds:
```

```
!!!!
```

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

```
!
```


```
CE2#ping 192.168.202.10
```

Type escape sequence to abort.
 Sending 5, 100-byte ICMP Echos to 192.168.202.10, timeout is 2 seconds:
 !!!!!
 Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms

Task 1: Configure BGP Weight and Local Preference

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:



Activity Procedure

Perform these steps:

- Step 1** 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 Gigabit Ethernet subinterface (configured in task 0) on the PE router.
- Step 2** On the pod CE router, use the per-neighbor BGP weight option to configure BGP weight. Configure the CE router to prefer BGP routes received from the neighbor pod PE router.
- Step 3** On the pod PE router, use 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.

Note The BGP routing protocol will not immediately update the routing table. Use Cisco IOS/IOS XE **clear ip bgp * soft** 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.

Activity Verification

You have completed this task when you attain these results:

- 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. The following outputs are taken from the pod CE routers:

```
CE1#show ip bgp
< text omitted>
*> 192.168.1.0      192.168.201.20      2 64500 i
*                  192.168.101.10      1 64500 i
*> 192.168.2.0      192.168.201.20      2 64500 i
*                  192.168.101.10      1 64500 i
!
CE2#show ip bgp
* 192.168.1.0      192.168.102.20      1 64500 i
*>                  192.168.202.10      2 64500 i
* 192.168.2.0      192.168.102.20      1 64500 i
*>                  192.168.202.10      2 64500 i
```

- The pod PE router should show a preference on the internal BGP route toward the pod CE router. The following outputs are taken from the PE (Cisco IOS XR and Cisco IOS XE) routers:


```
RP/0/RSP0/CPU0:PE1#show bgp
< text omitted >
*>i10.1.10.1/32    10.2.1.1            0    100    0 64501 i
*                  192.168.101.11      0    99     0 64501 i
!
PE2#show ip bgp
< text omitted >
*>i10.2.10.1/32    10.1.1.1            0    100    0 64502 i
*                  192.168.102.21      0    99     0 64502 i
```

Task 2: Configure BGP Multi-Exit-Discriminator

In this task, you will remove BGP local preference configuration from the previous task and configure the pod CE router to send BGP routes 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:

Visual Objective for Lab 5-2 Task 2: Configure BGP Multi-Exit-Discriminator



Activity Procedure

Complete these steps:

Step 1 From the pod PE routers, remove the local preference configuration created in the previous task.

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:PE1#show bgp
< text omitted>
*> 10.1.10.1/32      192.168.101.11      0          0 64501 i
!
PE2#show ip bgp
< text omitted>
* i10.2.10.1/32      10.1.1.1             0        100      0 64502 i
*>                  192.168.102.21       0          0 64502 i
```

Step 2 On the pod CE routers, use route maps to set 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.

Activity Verification

You have completed this task when you attain this result:


- Verify that the pod PE router selects the non-direct path to the pod CE router. The following outputs are taken from the PE routers.

```
RP/0/RSP0/CPU0:PE1#show bgp
< text omitted>
*>i10.1.10.1/32      10.2.1.1             1000      100      0 64501 i
*                  192.168.101.11       1001          0 64501 i
!
PE2#show ip bgp
< text omitted>
*>i10.2.10.1/32      10.1.1.1             1000      100      0 64502 i
*                  192.168.102.21       1001          0 64502 i
```

Task 3: Configure BGP Community

In this task, you will remove 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 service provider policy by sending BGP routes from the pod CE router with different BGP communities, and 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 applied to the BGP neighbors in the previous task.

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:PE1#show bgp
< text omitted >
*> 10.1.10.1/32      192.168.101.11      0          0 64501 i
!
PE2#show ip bgp
< text omitted >
* i10.2.10.1/32      10.1.1.1            0          100       0 64502 i
*>                  192.168.102.21      0          0 64502 i
```

- Step 2** On you team pod and the neighbor pod PE routers, use RPL (Cisco IOS XR) or **route-map** (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 3** On the pod and neighbor pod PE routers, apply RPL or route map to the BGP neighbors (pod CE router).
- Step 4** 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 5** 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 Backup_SP to the backup BGP neighbor (your pod PE router).

Activity Verification

You have completed this task when you attain these results:

- On the pod and neighbor pod PE routers, verify that BGP communities prepend to the BGP routes from the pod CE router:

```
RP/0/RSP0/CPU0:PE1#show bgp 10.1.10.1/32
Sat Nov 19 01:27:54.252 UTC
BGP routing table entry for 10.1.10.1/32
Versions:
  Process          bRIB/RIB   SendTblVer
  Speaker          349       349
Last Modified: Nov 19 01:27:01.861 for 00:00:52
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
  64501
    10.2.1.1 (metric 30) from 10.0.1.1 (10.2.1.1)
      Origin IGP, metric 0, localpref 101, valid, internal, best, group-best
      Received Path ID 0, Local Path ID 1, version 349
      Community: 64500:101
      Originator: 10.2.1.1, Cluster list: 10.0.1.1
  Path #2: Received by speaker 0
  Not advertised to any peer
  64501
    192.168.101.11 from 192.168.101.11 (11.0.0.1)
      Origin IGP, metric 0, localpref 99, valid, external
```

```

Received Path ID 0, Local Path ID 0, version 0
Community: 64500:99
!
PE2#show ip bgp 10.2.10.1/32
BGP routing table entry for 10.2.10.1/32, version 25
Paths: (2 available, best #1, table default)
  Advertised to update-groups:
    23
  64502
    10.1.1.1 (metric 30) from 10.0.1.1 (10.0.1.1)
      Origin IGP, metric 0, localpref 101, valid, internal, best
      Community: 64500:101
      Originator: 10.1.1.1, Cluster list: 10.0.1.1
  64502
    192.168.102.21 from 192.168.102.21 (12.0.0.1)
      Origin IGP, metric 0, localpref 99, valid, external
      Community: 64500:99

```

- On the pod and neighbor pod PE routers, verify the BGP routing table. Routes from the pod CE router have a different BGP local preference:

```

RP/0/RSP0/CPU0:PE1#show bgp
< text omitted >
*>i10.1.10.1/32      10.2.1.1      0      101      0 64501 i
*                   192.168.101.11    0      99      0 64501 i
!
PE2#show ip bgp
< text omitted >
*>i10.2.10.1/32      10.1.1.1      0      101      0 64502 i
*                   192.168.102.21    0      99      0 64502 i

```

- Step 6** On the pod CE router, modify the route-maps configured in the previous step. The primary_SP route-map should set BGP community 64500:999 and Backup_SP route-map should set BGP community 64500:1000.

- 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:

```

RP/0/RSP0/CPU0:PE1#show bgp 10.1.10.1/32
Sat Nov 19 01:30:38.738 UTC
BGP routing table entry for 10.1.10.1/32
Versions:
  Process          bRIB/RIB  SendTblVer
  Speaker          352      352
Last Modified: Nov 19 01:30:15.861 for 00:00:23
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
  64501
    10.2.1.1 (metric 30) from 10.0.1.1 (10.2.1.1)
      Origin IGP, metric 999, localpref 100, valid, internal, best, group-best
      Received Path ID 0, Local Path ID 1, version 352
      Community: 64500:999
      Originator: 10.2.1.1, Cluster list: 10.0.1.1
  Path #2: Received by speaker 0
  Not advertised to any peer
  64501
    192.168.101.11 from 192.168.101.11 (11.0.0.1)
      Origin IGP, metric 1000, localpref 100, valid, external
      Received Path ID 0, Local Path ID 0, version 0
      Community: 64500:1000
!
PE2#show ip bgp 10.2.10.1/32
BGP routing table entry for 10.2.10.1/32, version 26

```

```

Paths: (2 available, best #1, table default)
  Advertised to update-groups:
    23
  64502
    10.1.1.1 (metric 30) from 10.0.1.1 (10.0.1.1)
      Origin IGP, metric 999, localpref 100, valid, internal, best
      Community: 64500:999
      Originator: 10.1.1.1, Cluster list: 10.0.1.1
  64502
    192.168.102.21 from 192.168.102.21 (12.0.0.1)
      Origin IGP, metric 1000, localpref 100, valid, external
      Community: 64500:1000

```

- On the pod and neighbor pod PE routers, verify the BGP routing table. Routes from the pod CE router have different BGP MED (metric):

```

RP/0/RSP0/CPU0:PE1#show bgp
< text omitted >
*>i10.1.10.1/32      10.2.1.1          999      100      0 64501 i
*                   192.168.101.11   1000     0        0 64501 i
!
PE2#show ip bgp
< text omitted >
*>i10.2.10.1/32      10.1.1.1          999      100      0 64502 i
*                   192.168.102.21   1000     0        0 64502 i

```

- On the pod CE router, verify the route received from the neighbor pod CE router. Observe the BGP AS-path attribute:

```

CE1#show ip bgp
< text omitted >
*> 10.2.10.1/32      192.168.201.20      2 64500 64502 i
*                   192.168.101.10     1 64500 64502 i
!
CE2#show ip bgp
< text omitted >
* 10.1.10.1/32      192.168.102.20      1 64500 64501 i
*>                  192.168.202.10     2 64500 64501 i

```

- Step 7** On the pod CE router, modify the route maps configured in the previous step. The Primary_SP route map should set the BGP community 64500:1, and the Backup_SP route map should set the BGP community 64500:3.

Note 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, they have two routes to your CE router loopback interface. For example, from the CE1 **show ip bgp** output, there are two paths to the CE2 loopback interface. The best path selected is the one with the higher weight, even though 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.

- 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:

```

RP/0/RSP0/CPU0:PE1#show bgp 10.1.10.1/32
Mon Nov 21 09:16:49.611 UTC
BGP routing table entry for 10.1.10.1/32
Versions:
  Process          bRIB/RIB   SendTblVer
  Speaker          526       526
Last Modified: Nov 21 09:16:22.861 for 00:00:26
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
64501
 192.168.101.11 from 192.168.101.11 (11.0.0.1)
   Origin IGP, metric 0, localpref 100, valid, external, best, group-best
   Received Path ID 0, Local Path ID 1, version 526
   Community: 64500:1
!
PE2#show ip bgp 10.2.10.1/32
BGP routing table entry for 10.2.10.1/32, version 49
Paths: (2 available, best #2, table default)
  Advertised to update-groups:
    29          31
64502
 10.1.1.1 (metric 30) from 10.0.1.1 (10.0.1.1)
   Origin IGP, metric 0, localpref 100, valid, internal
   Community: 64500:1
   Originator: 10.1.1.1, Cluster list: 10.0.1.1
64502
 192.168.102.21 from 192.168.102.21 (12.0.0.1)
   Origin IGP, metric 0, localpref 100, valid, external, best
   Community: 64500:3

```

- On the pod and neighbor pod PE routers, verify the BGP routing table. Observe routes from the pod CE router:

```

RP/0/RSP0/CPU0:PE1#show bgp
< text omitted >
*> 10.1.10.1/32          192.168.101.11          0          0 64501 i
!
PE2#show ip bgp
< text omitted >
* i10.2.10.1/32          10.1.1.1          0          100          0 64502 i
*>          192.168.102.21          0          0 64502 i

```

- On the pod CE router, verify the routes received from the neighbor pod CE router. Observe that the BGP AS-path attribute is different for two routes:

```

CE1#show ip bgp
< text omitted >
*> 10.2.10.1/32          192.168.201.20          2 64500 64500 64500
64500 64502 i
*          192.168.101.10          1 64500 64500 64502
i
!
CE2#show ip bgp
< text omitted >
* 10.1.10.1/32          192.168.102.20          1 64500 64500 64500
64500 64501 i
*>          192.168.202.10          2 64500 64500 64501
i

```

Answer Key

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

Lab 2-1 Answer Key: Implement OSPF Routing

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

Task 1: Enable OSPF on the Routers

Step 1 The IP addresses configured on the CE and PE routers:

CE1 (Cisco IOS):

```
interface Loopback0
 ip address 10.1.10.1 255.255.255.255
!
interface GigabitEthernet0/0
 ip address 192.168.101.11 255.255.255.0
 no shutdown
```

PE1 (Cisco IOS XR):

```
interface Loopback0
 ipv4 address 10.1.1.1 255.255.255.255
!
interface GigabitEthernet0/0/0/0
 ipv4 address 192.168.101.10 255.255.255.0
 no shutdown
!
interface GigabitEthernet0/0/0/1
 ipv4 address 192.168.112.10 255.255.255.0
 no shutdown
!
interface GigabitEthernet0/0/0/2
 ipv4 address 192.168.11.10 255.255.255.0
 no shutdown
!
interface GigabitEthernet0/0/0/3
 ipv4 address 192.168.12.10 255.255.255.0
 no shutdown
!
commit
```

CE2 (Cisco IOS):

```
interface Loopback0
 ip address 10.2.10.1 255.255.255.255
!
interface GigabitEthernet0/0
 ip address 192.168.102.21 255.255.255.0
 no shutdown
```

PE2 (Cisco IOS XE):

```
interface Loopback0
 ip address 10.2.1.1 255.255.255.255
!
interface GigabitEthernet0/0/0
 ip address 192.168.102.20 255.255.255.0
 no shutdown
!
interface GigabitEthernet0/0/1
 ip address 192.168.112.20 255.255.255.0
 no shutdown
```

```

!
interface GigabitEthernet0/0/2
 ip address 192.168.21.20 255.255.255.0
 no shutdown
!
interface GigabitEthernet0/0/3
 ip address 192.168.22.20 255.255.255.0
 no shutdown

```

Step 3 Values used in the Team 1 for OSPF areas:

Link (Interface)	OSPF Area
PE1 (Gi0/0/0/2) – P1	0
PE1 (Gi0/0/0/3) – P2	0
PE2 (Gi0/0/0/2) – P1	0
PE2 (Gi0/0/0/3) – P2	0
PE1 (Gi0/0/0/1) – PE2 (Gi0/0/0/1)	0
PE1 (Gi0/0/0/0) – CE1 (Gi0/0)	11
PE2 (Gi0/0/0/0) – CE2 (Gi0/0)	12
PE1 (Loopback0)	1
CE1 (Loopback0)	11
PE2 (Loopback0)	2
CE2 (Loopback0)	12

Step 4 The OSPF configuration on the CE and PE routers for the Team 1:

CE1 (Cisco IOS):

```

router ospf 1
 network 10.1.10.1 0.0.0.0 area 11
 network 192.168.101.11 0.0.0.0 area 11

```

PE1 (Cisco IOS XR):

```

router ospf 1
 address-family ipv4
 area 0
 interface GigabitEthernet0/0/0/1
 interface GigabitEthernet0/0/0/2
 interface GigabitEthernet0/0/0/3
 area 1
 interface Loopback0
 area 11
 interface GigabitEthernet0/0/0/0
!
commit

```

CE2 (Cisco IOS):

```

router ospf 1
 network 10.2.10.1 0.0.0.0 area 12
 network 192.168.102.21 0.0.0.0 area 12

```

PE2 (Cisco IOS XE):

```

router ospf 1
 network 10.2.1.1 0.0.0.0 area 2
 network 192.168.21.20 0.0.0.0 area 0
 network 192.168.22.20 0.0.0.0 area 0
 network 192.168.102.20 0.0.0.0 area 12
 network 192.168.112.20 0.0.0.0 area 0

```

Step 6 The OSPF router IDs taken from Team 1:

Router	OSPF router ID
PE1 (Pod)	10.1.1.1
CE1 (Pod)	10.1.10.1
PE2 (neighbor Pod)	10.2.1.1
CE2 (neighbor Pod)	10.2.10.1

Step 7 Change of the OSPF router ID on the CE and PE routers for the Team 1:

CE1 (Cisco IOS):

```
router ospf 1
  router-id 10.10.10.10
end
!
clear ip ospf process
Reset ALL OSPF processes? [no]: yes
```

PE1 (Cisco IOS XR):

```
router ospf 1
  router-id 1.1.1.1
!
commit
!
clear ospf process
Reset ALL OSPF processes? [no]: yes
```

CE2 (Cisco IOS):

```
router ospf 1
  router-id 20.20.20.20
end
!
clear ip ospf process
Reset ALL OSPF processes? [no]: yes
```

PE2 (Cisco IOS XE):

```
router ospf 1
  router-id 2.2.2.2
end
!
clear ip ospf process
Reset ALL OSPF processes? [no]: yes
```

Step 8 The IPv6 enabled on the CE and PE routers:

CE1 (Cisco IOS):

```
ipv6 unicast-routing
!
interface Loopback0
  ipv6 address 2001:DB8:10:1:10::1/128
  ipv6 enable
!
interface GigabitEthernet0/0
  ipv6 enable
```

PE1 (Cisco IOS XR):

```
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
!
interface GigabitEthernet0/0/0/2
  ipv6 enable
!
interface GigabitEthernet0/0/0/3
  ipv6 enable
!
commit

```

CE2 (Cisco IOS):

```

ipv6 unicast-routing
!
interface Loopback0
  ipv6 address 2001:DB8:10:2:10::1/128
  ipv6 enable
!
interface GigabitEthernet0/0
  ipv6 enable

```

PE2 (Cisco IOS XE):

```

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
!
interface GigabitEthernet0/0/2
  ipv6 enable
!
interface GigabitEthernet0/0/3
  ipv6 enable

```

Step 9 The OSPFv3 enabled on the CE and PE routers:

CE1 (Cisco IOS):

```

interface Loopback0
  ipv6 ospf 1 area 11
!
interface GigabitEthernet0/0
  ipv6 ospf 1 area 11

```

PE1 (Cisco IOS XR):

```

router ospfv3 1
  area 0
    interface GigabitEthernet0/0/0/1
    !
    interface GigabitEthernet0/0/0/2
    !
    interface GigabitEthernet0/0/0/3
    !
  area 1
    interface Loopback0
    !
  area 11

```

```

    interface GigabitEthernet0/0/0/0
    !
commit

```

CE2 (Cisco IOS):

```

interface Loopback0
  ipv6 ospf 1 area 12
!
interface GigabitEthernet0/0
  ipv6 ospf 1 area 12

```

PE2 (Cisco IOS XE):

```

interface Loopback0
  ipv6 ospf 1 area 2
!
interface GigabitEthernet0/0/0
  ipv6 ospf 1 area 12
!
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

```

Task 2: Influence OSPF DR and BDR Election

Step 1 Configuration entered on the CE1 router:

```

interface GigabitEthernet0/0
  ip ospf priority 0

```

Configuration entered on the CE2 router:

```

interface GigabitEthernet0/0
  ip ospf priority 0

```

Step 2 Configuration entered on the PE1 router:

```

router ospf 1
  area 0
    interface GigabitEthernet0/0/0/2
      priority 2
    interface GigabitEthernet0/0/0/3
      priority 2
!
commit

```

Configuration entered on the PE2 router:

```

interface GigabitEthernet0/0/2
  ip ospf priority 2
interface GigabitEthernet0/0/3
  ip ospf priority 2

```

Task 3: Influence OSPF Route Selection by Changing OSPF Link Cost

Step 2 Configuration entered on the PE1 router:

```

router ospf 1
  area 0
    interface GigabitEthernet0/0/0/3
      cost 3
!
commit

```

Configuration entered on the PE2 router:

```
interface GigabitEthernet0/0/2
 ip ospf cost 3
```

Step 3 Change OSPF link cost between PE routers in your team.

Configuration entered on the PE1 router:

```
router ospf 1
 area 0
   interface GigabitEthernet0/0/0/1
     cost 10
 !
commit
```

Configuration entered on the PE2 router:

```
interface GigabitEthernet0/0/1
 ip ospf cost 10
```

Task 4: Configure OSPF Authentication

Step 1 The OSPF MD5 authentication is enabled on the CE and PE routers:

CE1 and CE2 (Cisco IOS):

```
interface GigabitEthernet0/0
 ip ospf authentication message-digest
 ip ospf message-digest-key 1 md5 cisco
```

PE1 (Cisco IOS XR):

```
router ospf 1
 area 11
   interface GigabitEthernet0/0/0/0
     authentication message-digest
     message-digest-key 1 md5 cisco
 !
commit
```

PE2 (Cisco IOS XE):

```
interface GigabitEthernet0/0/0
 ip ospf authentication message-digest
 ip ospf message-digest-key 1 md5 cisco
```

Task 5: Configure OSPF Virtual Links

Step 1 Enable new Loopback interface on the CE router:

CE1 (Cisco IOS):

```
interface Loopback11
 ip address 11.0.0.1 255.255.255.0
```

CE2 (Cisco IOS):

```
interface Loopback12
 ip address 12.0.0.1 255.255.255.0
```

Step 2 Add network into OSPF on the CE router:

CE1 (Cisco IOS):

```
router ospf 1
 network 11.0.0.1 0.0.0.0 area 100
```

CE2 (Cisco IOS):

```
router ospf 1
 network 12.0.0.1 0.0.0.0 area 200
```

Step 3 The OSPF virtual link is enabled between the CE and PE routers:

CE1 (Cisco IOS):

```
router ospf 1
 area 11 virtual-link 1.1.1.1
```

PE1 (Cisco IOS XR):

```
router ospf 1
 area 11
  virtual-link 10.10.10.10
!
commit
```

CE2 (Cisco IOS):

```
router ospf 1
 area 12 virtual-link 2.2.2.2
```

PE2 (Cisco IOS XE):

```
router ospf 1
 area 12 virtual-link 20.20.20.20
```

Lab 2-2 Answer Key: Implement OSPF Special Area Types

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

Task 1: Enable OSPF Summarization

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

CE1 (Cisco IOS):

```
router ospf 1
 no area 11 virtual-link 1.1.1.1
 no network 11.0.0.1 0.0.0.0 area 100
```

PE1 (Cisco IOS XR):

```
router ospf 1
 area 11
  no virtual-link 10.10.10.10
!
commit
```

CE2 (Cisco IOS):

```
router ospf 1
 no area 12 virtual-link 2.2.2.2
 no network 12.0.0.1 0.0.0.0 area 200
```

PE2 (Cisco IOS XE):

```
router ospf 1
 no area 12 virtual-link 20.20.20.20
```

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

CE1 (Cisco IOS):

```
router ospf 1
```

```
redistribute connected subnets route-map Loopback11
!
route-map Loopback11 permit 10
match interface Loopback11
```

CE2 (Cisco IOS):

```
router ospf 1
redistribute connected subnets route-map Loopback12
!
route-map Loopback12 permit 10
match interface Loopback12
```

Step 3 Configure OSPF summarization on the PE router:

PE1 (Cisco IOS XR):

```
router ospf 1
area 11
range 10.1.10.0/24
!
commit
```

PE2 (Cisco IOS XE):

```
router ospf 1
area 12 range 10.2.10.0 255.255.255.0
```

Step 4 Configure OSPF summarization on the CE router:

CE1 (Cisco IOS):

```
router ospf 1
summary-address 11.0.0.0 255.0.0.0
```

CE2 (Cisco IOS):

```
router ospf 1
summary-address 12.0.0.0 255.0.0.0
```

Task 2: Configure OSPF Stub Area

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

CE1 (Cisco IOS):

```
router ospf 1
area 11 stub
```

PE1 (Cisco IOS XR):

```
router ospf 1
area 11
stub
!
commit
```

CE2 (Cisco IOS):

```
router ospf 1
area 12 stub
```

PE2 (Cisco IOS XE):

```
router ospf 1
area 12 stub
```

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

PE1 (Cisco IOS XR):

```
router ospf 1
 area 11
  stub no-summary
!
commit
```

PE2 (Cisco IOS XE):

```
router ospf 1
 area 12 stub no-summary
```

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

PE1 (Cisco IOS XR):

```
router ospf 1
 area 11
  default-cost 10
!
commit
```

PE2 (Cisco IOS XE):

```
router ospf 1
 area 12 default-cost 10
```

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

CE1 (Cisco IOS):

```
router ospf 1
 no area 11 stub
 area 11 nssa
```

PE1 (Cisco IOS XR):

```
router ospf 1
 area 11
  no stub
  nssa no-summary
!
commit
```

CE2 (Cisco IOS):

```
router ospf 1
 no area 12 stub
 area 12 nssa
```

PE2 (Cisco IOS XE):

```
router ospf 1
 no area 12 stub
 area 12 nssa no-summary
```

Lab 3-1 Answer Key: Implement Integrated IS-IS Routing

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

Task 1: Enable Integrated IS-IS on the Routers

Step 3 Values used in Team 1 for IS-IS NET addressing, and interfaces used in IS-IS routing, are found in this table:

Router	IS-IS Area	System ID+00	IS-IS Interfaces
CE1	49.0001	0100.0101.0001.00	Gi0/0, Lo0
PE1	49.0001	0100.0100.1001.00	Gi0/0/0/0, Gi0/0/0/1, Lo0
CE2	49.0002	0100.0201.0001.00	Gi0/0, Lo0
PE2	49.0002	0100.0200.1001.00	Gi0/0/0, Gi0/0/1, Lo0

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

CE1 (Cisco IOS):

```
no router ospf 1
!
interface Loopback0
 ip router isis
!
interface GigabitEthernet0/0
 ip router isis
!
router isis
 net 49.0001.0100.0101.0001.00
```

PE1 (Cisco IOS XR):

```
no router ospf 1
!
router isis 1
 net 49.0001.0100.0100.1001.00
 address-family ipv6 unicast
  single-topology
!
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
```

CE2 (Cisco IOS):

```
no router ospf 1
!
interface Loopback0
 ip router isis
!
interface GigabitEthernet0/0
 ip router isis
!
router isis
 net 49.0002.0100.0201.0001.00
```

PE2 (Cisco IOS XE):

```
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 49.0002.0100.0200.1001.00
```

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

CE1 and CE2 (Cisco IOS):

```
router isis
 is-type level-1
```

PE1 (Cisco IOS XR):

```
router isis 1
 interface GigabitEthernet0/0/0/0
  circuit-type level-1
!
 interface GigabitEthernet0/0/0/1
  circuit-type level-2-only
!
commit
```

PE2 (Cisco IOS XE):

```
interface GigabitEthernet0/0/0
 isis circuit-type level-1
!
interface GigabitEthernet0/0/1
 isis circuit-type level-2-only
```

Task 2: IS-IS Route Summarization

Step 3 Configuration entered on the PE1 router:

```
router isis 1
 address-family ipv4 unicast
  summary-prefix 10.1.10.0/24 level 2
!
commit
```

Configuration entered on the PE2 router:

```
router isis
 summary-address 10.2.10.0 255.255.255.0
```

Task 3: Enable IPv6 IS-IS Routing

Step 4 The IPv6 IS-IS configuration on the CE and PE routers:

CE1 and CE2(Cisco IOS):

```
no ipv6 router ospf 1
!
interface Loopback0
 ipv6 router isis
!
interface GigabitEthernet0/0
 ipv6 router isis
```

PE1 (Cisco IOS XR):

```
no router ospfv3 1
!
router isis 1

!
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
```

PE2 (Cisco IOS XE):

```
no ipv6 router ospf 1
!
interface Loopback0
 ipv6 router isis
!
interface GigabitEthernet0/0/0
 ipv6 router isis
!
interface GigabitEthernet0/0/1
 ipv6 router isis
```

Lab 4-1 Answer Key: Implement Basic BGP Routing

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

Task 1: Configure External BGP

Step 1 The BGP configuration on the PE routers:

PE1 (Cisco IOS XR):

```
router bgp 64500
 address-family ipv4 unicast
!
 neighbor 192.168.101.11
  remote-as 64501
 address-family ipv4 unicast
!
commit
```

PE2 (Cisco IOS XE):

```
router bgp 64500
 neighbor 192.168.102.21 remote-as 64502
```

Step 2 The BGP configuration on the CE routers:

CE1 (Cisco IOS):

```
router bgp 64501
 neighbor 192.168.101.10 remote-as 64500
```

CE2 (Cisco IOS):

```
router bgp 64502
```

```
neighbor 192.168.102.20 remote-as 64500
```

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

CE1 (Cisco IOS):

```
router bgp 64501
neighbor 192.168.101.10 password cisco
```

PE1 (Cisco IOS XR):

```
router bgp 64500
neighbor 192.168.101.11
password clear cisco
!
commit
```

CE2 (Cisco IOS):

```
router bgp 64502
neighbor 192.168.102.20 password cisco
```

PE2 (Cisco IOS XE):

```
router bgp 64500
neighbor 192.168.102.21 password cisco
```

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

CE1 (Cisco IOS):

```
router bgp 64501
network 10.1.10.1 mask 255.255.255.255
```

PE1 (Cisco IOS XR):

```
router bgp 64500
address-family ipv4 unicast
network 10.1.1.1/32
!
commit
```

CE2 (Cisco IOS):

```
router bgp 64502
network 10.2.10.1 mask 255.255.255.255
```

PE2 (Cisco IOS XE):

```
router bgp 64500
network 10.2.1.1 mask 255.255.255.255
```

Step 6 Configure and apply inbound and outbound BGP route policy on the PE (Cisco IOS XR only) router:

PE1 (Cisco IOS XR):

```
route-policy pass
pass
end-policy
!
router bgp 64500
neighbor 192.168.101.11
address-family ipv4 unicast
route-policy pass in
route-policy pass out
!
commit
```

Task 2: Configure Internal BGP

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

PE1 (Cisco IOS XR):

```
interface GigabitEthernet0/0/0/2
  cdp
  ipv4 address 192.168.11.10 255.255.255.0
  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
```

PE2 (Cisco IOS XE):

```
interface GigabitEthernet0/0/2
  ip address 192.168.21.20 255.255.255.0
  ip router isis
  ipv6 enable
  ipv6 router isis
  cdp enable
  isis circuit-type level-2-only
```

Step 2 The BGP configuration on the PE routers:

PE1 (Cisco IOS XR):

```
router bgp 64500
  neighbor 10.0.1.1
  remote-as 64500
  update-source Loopback0
  address-family ipv4 unicast
!
commit
```

PE2 (Cisco IOS XE):

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

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

PE1 (Cisco IOS XR):

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

PE2 (Cisco IOS XE):

```
router bgp 64500
  neighbor 10.0.1.1 next-hop-self
```

Lab 5-1 Answer Key: Implement Route Redistribution

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

Task 1: Redistribute Between OSPF and IS-IS

Step 3 Disable IS-IS and start OSPF on the CE routers:

CE1 (Cisco IOS):

```
no router isis
router ospf 1
 network 192.168.101.11 0.0.0.0 area 11
 network 10.1.10.1 0.0.0.0 area 11
```

CE2 (Cisco IOS):

```
no router isis
router ospf 1
 network 192.168.102.21 0.0.0.0 area 12
 network 10.2.10.1 0.0.0.0 area 12
```

Step 4 Disable IS-IS and start OSPF on the PE routers:

PE1 (Cisco IOS XR):

```
router isis 1
 no interface GigabitEthernet0/0/0/0
 no interface GigabitEthernet0/0/0/1
!
router ospf 1
 area 11
  interface GigabitEthernet0/0/0/0
   authentication message-digest
   message-digest-key 1 md5 cisco
 area 0
  interface GigabitEthernet0/0/0/1
!
commit
```

PE2 (Cisco IOS XE):

```
interface GigabitEthernet0/0/0
 no ip router isis
 no ipv6 router isis
 no isis circuit-type level-1
!
interface GigabitEthernet0/0/1
 no ip router isis
 no ipv6 router isis
 no isis circuit-type level-2-only
!
router ospf 1
 network 192.168.102.20 0.0.0.0 area 12
 network 192.168.112.20 0.0.0.0 area 0
```

Step 4 Redistribute between IS-IS and OSPF on the pod PE routers:

PE1 (Cisco IOS XR):

```
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
```

PE2 (Cisco IOS XE):

```
router ospf 1
```

```

    redistribute isis level-2 subnets
!
router isis
    redistribute ospf 1 level-2

```

Task 2: Prevent Potential Routing Loops

Step 1 Configure RPL or route map on the PE routers:

PE1 (Cisco IOS XR):

```

route-policy OSPFintoISIS
    if tag eq 115 then drop else set tag 110 endif
exit
!
commit

```

PE2 (Cisco IOS XE):

```

route-map OSPFintoISIS deny 10
    match tag 115
!
route-map OSPFintoISIS permit 20
    set tag 110

```

Step 2 Apply RPL or route map to the OSPF into IS-IS redistribution:

PE1 (Cisco IOS XR):

```

router isis 1
    address-family ipv4 unicast
        no redistribute ospf 1 level-2
        redistribute ospf 1 level-2 route-policy OSPFintoISIS
!
commit

```

PE2 (Cisco IOS XE):

```

router isis
    no redistribute ospf 1
    redistribute ospf 1 route-map OSPFintoISIS

```

Step 3 Configure RPL or route map on the PE routers:

PE1 (Cisco IOS XR):

```

route-policy ISISintoOSPF
    if tag eq 110 then drop else set tag 115 endif
exit
!
commit

```

PE2 (Cisco IOS XE):

```

route-map ISISintoOSPF deny 10
    match tag 110
!
route-map ISISintoOSPF permit 20
    set tag 115

```

Step 4 Apply RPL or route map to the IS-IS into OSPF redistribution:

PE1 (Cisco IOS XR):

```

router ospf 1
    no redistribute isis 1 level-2
    redistribute isis 1 level-2 route-policy ISISintoOSPF
!
commit

```

PE2 (Cisco IOS XE):

```
router ospf 1
  no redistribute isis
  redistribute isis level-2 subnets route-map ISISintoOSPF
```

Task 3: Modify Administrative Distance

Step 1 Modify administrative distance on the PE routers:

PE1 (Cisco IOS XR):

```
router ospf 1
  distance ospf external 116
!
commit
```

PE2 (Cisco IOS XE):

```
router ospf 1
  distance ospf external 116
```

Task 4: One-way Redistribution

Step 1 Remove IS-IS into OSPF redistribution on the PE routers:

PE1 (Cisco IOS XR):

```
router ospf 1
  no redistribute isis 1
!
commit
```

PE2 (Cisco IOS XE):

```
router ospf 1
  no redistribute isis
```

Step 2 Originate OSPF default route on the PE routers:

PE1 (Cisco IOS XR):

```
router ospf 1
  default-information originate always
!
commit
```

PE2 (Cisco IOS XE):

```
router ospf 1
  default-information originate always
```

Lab 5-2 Answer Key: Influence BGP Route Selection

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

Task 0: Set Second Link Between CE and PE

Step 1 Enable second link between CE and PE routers:

SW1 (Cisco IOS):

```
vlan 20
!
interface FastEthernet0/2
  port-type nni
  switchport mode trunk
  no shutdown
!
```

```
interface FastEthernet0/23
  port-type nni
  switchport access vlan 20
  no shutdown
```

SW2 (Cisco IOS):

```
vlan 10
!
interface FastEthernet0/2
  port-type nni
  switchport mode trunk
  no shutdown
!
interface FastEthernet0/23
  port-type nni
  switchport access vlan 10
  no shutdown
```

SW12 (Cisco IOS):

```
vlan 10,20
!
interface FastEthernet0/1
  port-type nni
  switchport access vlan 10
  no shutdown
!
interface FastEthernet0/2
  port-type nni
  switchport access vlan 20
  no shutdown
!
interface FastEthernet0/21
  port-type nni
  switchport access vlan 20
  no shutdown
!
interface FastEthernet0/23
  port-type nni
  switchport access vlan 10
  no shutdown
```

CE1 (Cisco IOS):

```
interface GigabitEthernet0/1
  ip address 192.168.201.11 255.255.255.0
  no shutdown
  speed 100
```

PE1 (Cisco IOS XR):

```
interface GigabitEthernet0/0/0/0.2
  ipv4 address 192.168.202.10 255.255.255.0
  encapsulation dot1q 20
!
commit
```

CE2 (Cisco IOS):

```
interface GigabitEthernet0/1
  ip address 192.168.202.21 255.255.255.0
  no shutdown
  speed 100
```

PE2 (Cisco IOS XE):

```
interface GigabitEthernet0/0/0.1
  encapsulation dot1q 10
```

```
ip address 192.168.201.20 255.255.255.0
```

Step 2 Disable OSPF between CE and PE routers:

CE1 (Cisco IOS):

```
no router ospf 1
```

PE1 (Cisco IOS XR):

```
no router ospf 1
!
commit
```

CE2 (Cisco IOS):

```
no router ospf 1
```

PE2 (Cisco IOS XE):

```
no router ospf 1
```

Task 1: Configure BGP Weight and Local Preference

Step 1 Enable second EBGP session between CE and PE routers:

CE1 (Cisco IOS):

```
router bgp 64501
 neighbor 192.168.201.20 remote-as 64500
```

PE1 (Cisco IOS XR):

```
router bgp 64500
 neighbor 192.168.202.21
   remote-as 64502
   address-family ipv4 unicast
     route-policy pass out
     route-policy pass in
!
commit
```

CE2 (Cisco IOS):

```
router bgp 64502
 neighbor 192.168.202.10 remote-as 64500
```

PE2 (Cisco IOS XE):

```
router bgp 64500
 neighbor 192.168.201.11 remote-as 64501
```

Step 2 Configure BGP weight:

CE1 (Cisco IOS):

```
router bgp 64501
 neighbor 192.168.101.10 weight 1
 neighbor 192.168.201.20 weight 2
```

CE2 (Cisco IOS):

```
router bgp 64502
 neighbor 192.168.102.20 weight 1
 neighbor 192.168.202.10 weight 2
```

Step 3 Configure BGP local preference:

PE1 (Cisco IOS XR):

```
route-policy LP99
```

```

    set local-preference 99
end-policy
!
router bgp 64500
  neighbor 192.168.101.11
    address-family ipv4 unicast
      route-policy LP99 in
    !
  commit

```

PE2 (Cisco IOS XE):

```

route-map LP99 permit
  set local-preference 99
!
router bgp 64500
  neighbor 192.168.102.21 route-map LP99 in

```

Task 2: Configure BGP Multi-Exit-Discriminator

Step 1 Remove BGP local preference configuration:

PE1 (Cisco IOS XR):

```

router bgp 64500
  neighbor 192.168.101.11
    address-family ipv4 unicast
      no route-policy LP99 in
      route-policy pass in
    !
  commit

```

PE2 (Cisco IOS XE):

```

router bgp 64500
  no neighbor 192.168.102.21 route-map LP99 in

```

Step 2 Send BGP MED to the BGP neighbors:

CE1 (Cisco IOS):

```

route-map MED1000 permit
  set metric 1000
!
route-map MED1001 permit
  set metric 1001
!
router bgp 64501
  neighbor 192.168.101.10 route-map MED1001 out
  neighbor 192.168.201.20 route-map MED1000 out

```

CE2 (Cisco IOS):

```

route-map MED1000 permit
  set metric 1000
!
route-map MED1001 permit
  set metric 1001
!
router bgp 64502
  neighbor 192.168.102.20 route-map MED1001 out
  neighbor 192.168.202.10 route-map MED1000 out

```

Task 3: Configure BGP Community

Step 1 Remove route map:

CE1 (Cisco IOS):

```

router bgp 64501
  no neighbor 192.168.101.10 route-map MED1001 out

```

```
no neighbor 192.168.201.20 route-map MED1000 out
```

CE2 (Cisco IOS):

```
router bgp 64502
no neighbor 192.168.102.20 route-map MED1001 out
no neighbor 192.168.202.10 route-map MED1000 out
```

Step 2 Configure BGP community policy:

PE1 (Cisco IOS XR):

```
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
```

PE2 (Cisco IOS XE):

```
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

```

Step 3 Apply RPL and route map:

PE1 (Cisco IOS XR):

```

router bgp 64500
neighbor 192.168.101.11
address-family ipv4 unicast
route-policy to_Cust out
route-policy from_Cust in
neighbor 192.168.202.21
address-family ipv4 unicast
route-policy to_Cust out
route-policy from_Cust in
!
commit

```

PE2 (Cisco IOS XE):

```

router bgp 64500
neighbor 192.168.201.11 route-map to_Cust out
neighbor 192.168.102.21 route-map to_Cust out
neighbor 192.168.201.11 route-map from_Cust in
neighbor 192.168.102.21 route-map from_Cust in

```

Step 4 Enable BGP community support:

CE1 (Cisco IOS):

```

ip bgp-community new-format
!
router bgp 64501
neighbor 192.168.101.10 send-community both
neighbor 192.168.201.20 send-community both

```

CE2 (Cisco IOS):

```

ip bgp-community new-format
!
router bgp 64502
neighbor 192.168.102.20 send-community both
neighbor 192.168.202.10 send-community both

```

PE2 (Cisco IOS XE):

```

ip bgp-community new-format
!
router bgp 64500
neighbor 10.0.1.1 send-community both

```

Step 5 Set BGP communities to change BGP local preference:

CE1 (Cisco IOS):

```
route-map Primary_SP permit 10
  set community 64500:101
route-map Backup_SP permit 10
  set community 64500:99
!
router bgp 64501
  neighbor 192.168.101.10 route-map Backup_SP out
  neighbor 192.168.201.20 route-map Primary_SP out
```

CE2 (Cisco IOS):

```
route-map Primary_SP permit 10
  set community 64500:101
route-map Backup_SP permit 10
  set community 64500:99
!
router bgp 64502
  neighbor 192.168.102.20 route-map Backup_SP out
  neighbor 192.168.202.10 route-map Primary_SP out
```

Step 6 Set BGP communities to change BGP MED:

CE1 (Cisco IOS):

```
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
```

CE2 (Cisco IOS):

```
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
```

Step 7 Set BGP communities to change BGP AS-path prepend:

CE1 (Cisco IOS):


```
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
```

CE2 (Cisco IOS):


```
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
```

Appendix A: Lab Topology

SPROUTe Lab Topology




SPROUTe Lab Interface Identification: Team View



© 2012 Cisco and/or its affiliates. All rights reserved. SPROUTE v1.01-5


SPROUTe Lab IP Addressing: Team View



© 2012 Cisco and/or its affiliates. All rights reserved.

SPROUTe v1.01-6

SPROUTe Management IP Addresses



SPROUTe v1.01-7

© 2012 Cisco and/or its affiliates. All rights reserved.

