



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

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

Europe Headquarters
Cisco Systems International BV
Amsterdam,
The Netherlands

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

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

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

Copyright Date:

© 2014 Cisco Systems, Inc.

Table of Contents

Lab Introduction	1
Lab Topology Diagram (Backbone Pod View)	1
Lab Topology Diagram (Student Pod View)	2
Global Job Aids	3
Pod Access Information	3
Device Information	4
IP Addressing	6
Hardware Lab 1: Implement BGP Route Reflectors	9
Visual Objective	10
Command List	10
Task 1: Verify the Existing BGP Sessions	13
Task 2: Configure a Route Reflector and Internal BGP Session	18
Task 3: Restrict Route Propagation to a Client	22
Hardware Lab 2: Implement BGP Security Options	25
Visual Objective	26
Command List	26
Task 1: Implement BGP Neighbor Authentication Using Passwords	29
Task 2: Implement BGP Neighbor Authentication Using Key Chains	32
Task 3: Enable BGP TTL Security Check	34
Task 4: Enable CoPP	36
Task 5: Enable RTBH Filtering	38
Hardware Lab 3: Improve BGP Scalability	43
Visual Objective	44
Command List	44
Task 1: Implement the BGP Configuration and Peer Templates	47
Task 2: Limit the Number of Prefixes Received from a BGP Neighbor	53
Task 3: Improve BGP Convergence by Changing the BGP Scan and Advertisement Interval	55
Task 4: Implement BGP Route Dampening	61
Hardware Lab 4: Implement Layer 2 and Layer 3 Multicast	67
Visual Objective	68
Command List	68
Task 1: Enable IGMP and MLD	70
Task 2: Verify IGMP Snooping	73
Hardware Lab 5: Enable and Optimize PIM-SM	75
Visual Objective	76
Command List	76
Task 1: Implement PIM-SM	78
Task 2: Shared Tree Formation—Receivers	81
Task 3: Shared Tree Formation—Sources	85
Task 4: Switching to the SPT	92
Hardware Lab 6: Implement PIM-SM Enhancements	99
Visual Objective	100
Command List	100
Task 1: Implement PIM-SSM	102
Hardware Lab 7: Implement Rendezvous Point Distribution	109
Visual Objective	110
Command List	110
Task 1: Enable OSPF Routing	112


Task 2: Enable BSR.....	117
Task 3: Enable Anycast RP	120
Hardware Lab 8: Implement a DHCPv6 Server with Prefix Delegation.....	125
Visual Objective.....	126
Command List.....	126
Task 1: Configure a Prefix Delegation DHCPv6 Server and Client.....	127
Task 2: Configure DHCPv6 Lite Server.....	129
Hardware Lab 9: Implement IPv6 Multicasting.....	131
Visual Objective.....	132
Command List.....	132
Task 1: Create a New Loopback Interface and Verify Connectivity	134
Task 2: Implement IPv6 Multicast Using Embedded RPs	137
Hardware Lab 10: Implement Tunnels for IPv6	145
Visual Objective.....	146
Command List.....	146
Task 1: Configure a Static IPv6-in-IPv4 Tunnel.....	147
Task 2: Configure Dynamic 6RD Tunnels.....	150
Answer Key	157
Hardware Lab 1: Implement BGP Route Reflectors	157
Task 2: Configure a Route Reflector and Internal BGP Session.....	157
Task 3: Restrict Route Propagation to a Client.....	159
Hardware Lab 2: Implement BGP Security Options	160
Task 1: Implement BGP Neighbor Authentication Using Passwords.....	160
Task 2: Implement BGP Neighbor Authentication Using Key Chains	161
Task 3: Enable BGP TTL Security Check	162
Task 4: Enable CoPP	163
Task 5: Enable RTBH Filtering.....	164
Hardware Lab 3: Improve BGP Scalability	167
Task 1: Implement the BGP Configuration and Peer Templates	167
Task 2: Limit the Number of Prefixes Received from a BGP Neighbor	168
Task 3: Improve BGP Convergence by Changing the BGP Scan and Advertisement Interval	168
Task 4: Implement BGP Route Dampening.....	169
Hardware Lab 4: Implement Layer 2 and Layer 3 Multicast.....	171
Task 1: Enable IGMP and MLD.....	171
Task 2: Verify IGMP Snooping	172
Hardware Lab 5: Enable and Optimize PIM-SM.....	174
Task 1: Implement PIM-SM	174
Task 2: Shared Tree Formation—Receivers	176
Task 3: Shared Tree Formation—Sources	178
Task 4: Switching to the SPT.....	179
Hardware Lab 6: Implement PIM-SM Enhancements.....	181
Task 1: Implement PIM-SSM	181
Hardware Lab 7: Implement Rendezvous Point Distribution.....	184
Task 1: Enable OSPF Routing	184
Task 2: Enable BSR.....	186
Task 3: Enable Anycast RP	187
Hardware Lab 8: Implement a DHCPv6 Server with Prefix Delegation.....	190
Task 1: Configure a Prefix Delegation DHCPv6 Server and Client.....	190
Task 2: Configure DHCPv6 Lite Server.....	190
Hardware Lab 9: Implement IPv6 Multicasting	192
Task 1: Create a New Loopback Interface and Verify Connectivity	192
Task 2: Implement IPv6 Multicast Using Embedded RPs	193
Hardware Lab 10: Implement Tunnels for IPv6.....	195

Task 1: Configure a Static IPv6-in-IPv4 Tunnel	195
Task 2: Configure Dynamic 6RD Tunnels.....	196


Lab Introduction

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

Lab Topology Diagram (Backbone Pod View)



Lab Topology Diagram (Student Pod View)



Global Job Aids

Pod Access Information

The instructor will provide the information for all team and pod number access. Write down the information in the table for future reference.


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

Device Information

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

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

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




Device Roles and Loopback IP Addresses

Device Name	Device Role	Lo0 IPv4 Address	Lo0 IPv6 Address
CEx 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


Device Name	Device Role	Lo0 IPv4 Address	Lo0 IPv6 Address
P1	Cisco ASR 9000 core router	10.0.1.1/32	2001:db8:10:0:1::1/128
P2	Cisco ASR 9000 core router	10.0.2.1/32	2001:db8:10:0:2::1/128

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




IP Addressing

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



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



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

Pod IP Addressing

Device	Interface	IPv4 Address	IPv6 Address
CEx	Gi0/0	192.168.10x.x1/24	2001:db8:192:168:10x::x1/80
CEy	Gi0/0	192.168.10y.y1/24	2001:db8:192:168:10y::y1/80
P1		192.168.x1.1/24	2001:db8:192:168:x1::1/80
		192.168.y1.1/24	2001:db8:192:168:y1::1/80
P2		192.168.x2.2/24	2001:db8:192:168:x2::2/80
		192.168.y2.2/24	2001:db8:192:168:y2::2/80
PE2	POS0/2/0	192.168.211.20/24	2001:db8:192:168:211::20/80
	POS0/2/1	192.168.212.20/24	2001:db8:192:168:212::20/80
PE4	POS0/2/0	192.168.211.40/24	2001:db8:192:168:211::40/80
	POS0/2/1	192.168.212.40/24	2001:db8:192:168:212::40/80
PEx	Gi0/0/0/0	192.168.10x.x0/24	2001:db8:192:168:10x::x0/80
	Gi0/0/0/1	192.168.1xy.x0/24	2001:db8:192:168:1xy::x0/80
	Gi0/0/0/2	192.168.x1.x0/24	2001:db8:192:168:x1::x0/80
	Gi0/0/0/3	192.168.x2.x0/24	2001:db8:192:168:x2::x0/80
PEy	Gi0/0/0	192.168.10y.y0/24	2001:db8:192:168:10y::y0/80
	Gi0/0/1	192.168.1xy.y0/24	2001:db8:192:168:1xy::y0/80
	Gi0/0/2	192.168.y1.y0/24	2001:db8:192:168:y1::y0/80
	Gi0/0/3	192.168.y2.y0/24	2001:db8:192:168:y2::y0/80

Core IP Addressing

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

Hardware Lab 1: Implement BGP Route Reflectors

Activity Objective

In this lab activity, you will configure a backbone router (P2) to act as a redundant route reflector. You will also configure an additional IBGP session between the PE router in your pod and the redundant route reflector (P2). Your instructor already configured P1 as a route reflector for the pod PE routers.


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

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

- Verify existing EBGP and IBGP sessions
- Configure a route reflector and IBGP session between a pod router and backbone router
- Restrict route propagation to a route reflector client

Visual Objective

The figure illustrates what you will accomplish in this activity.



Command List

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

Cisco IOS/IOS XE Commands

Command	Description
[no] shutdown	Enables or disables the interface on the router.
configure terminal	Enters configuration mode.
interface <i>interface</i>	Enters interface configuration mode.
ip address <i>ip_address subnet_mask</i>	Sets a primary or secondary IPv4 address for an interface and the subnet mask.
[ip ipv6] router <i>isis</i>	Enables IS-IS routing to the interface.
ipv6 address <i>ip_address/prefix</i>	Sets the IPv6 address for an interface.
ipv6 enable	Enables IPv6 support on the interface.

Command	Description
isis circuit-type [level-1 level-1-2 level-2-only]	Enables the IS-IS routing process to establish a selected IS-IS circuit type on the interface.
neighbor <i>IP_address</i> next-hop-self	Enables the internal BGP router to send BGP routes with the router BGP IP address.
neighbor <i>IP_address</i> remote-as <i>AS-number</i>	Configures the BGP neighbor.
neighbor <i>IP_address</i> update-source <i>interface</i>	Enables the BGP router to send BGP packets from the source IP address.
ping <i>dest_IP</i> <i>source</i> <i>source_IP</i>	Verifies connectivity between the source IP and destination IP.
router bgp <i>AS-number</i>	Creates a BGP process and enters BGP process configuration mode.
show ip bgp [<i>prefix</i>]	Displays the BGP routing table.
show ip bgp summary	Displays the BGP routing protocol characteristics, including BGP neighbor status.
show ip interface brief	Displays the interface status and the IPv4 addresses that are configured.
show isis neighbors	Displays the IS-IS neighbor information.

Cisco IOS XR Commands

Command	Description
[no] shutdown	Enables or disables the interface on the router.
address-family [ipv4 ipv6] unicast	Enables IPv4 or IPv6 IS-IS or BGP routing and enters address family configuration mode for IS-IS or BGP (in router IS-IS or BGP configuration mode).
bgp cluster-id <i>cluster_ID</i>	Sets the route reflector cluster ID.
circuit-type [level-1 level-1-2 level-2-only]	Enables the IS-IS routing process to establish the selected IS-IS circuit type on the interface.
commit	Commits changes to the running configuration.
configure terminal	Enters configuration mode.
interface <i>interface</i> (global command)	Enters interface configuration mode.
interface <i>interface</i> (router command)	Defines the interfaces on which the IS-IS protocol runs.

Command	Description
ip address <i>ip_address mask</i>	Sets the IPv4 address for an interface.
ipv6 address <i>ip_address/prefix</i>	Sets the IPv6 address for an interface.
neighbor <i>IP-address</i>	Configures the BGP neighbor and enters BGP neighbor configuration mode.
next-hop-self	Enables the internal BGP router to send BGP routes with the router BGP IP address (BGP neighbor address family mode).
pass	Passes the route for further processing (route-policy configuration mode).
ping <i>dest_IP source source_IP</i>	Verifies connectivity between the source IP and destination IP (IPv4 and IPv6).
remote-as <i>AS_number</i>	Configures the AS number for the BGP neighbor (BGP neighbor mode).
route-policy <i>route_policy_name</i>	Creates the route policy and enters route policy configuration mode.
route-policy <i>route_policy_name</i> [in out]	Applies the route policy to the BGP neighbor.
router bgp <i>AS_number</i>	Creates a BGP process and enters BGP process configuration mode.
router isis <i>process_ID</i>	Creates an IS-IS process.
route-reflector-client	Configures an IBGP neighbor as the route reflector client.
show bgp [<i>prefix</i>]	Displays the BGP routing table.
show bgp summary	Displays the BGP routing protocol characteristics, including the BGP neighbor status.
show ipv4 interface brief	Displays the interface status and the IPv4 addresses that are configured.
show isis neighbors	Displays the IS-IS neighbor information.
update-source <i>interface</i>	Enables the BGP router to send the BGP packets from the source IP address (BGP neighbor configuration mode).

Task 1: Verify the Existing BGP Sessions

In this task, you will verify that BGP is already configured and running in your pod, as preconfigured by your instructor. An IBGP session should be established between the PE router in your pod and the P1 backbone router. An EBGP session should be established between the PE and CE routers.

Activity Procedure

Complete these steps:

- Step 1** Use Telnet to connect to the P1 router. Examine the BGP configuration. You should see that the PE routers are configured as route reflector clients.

```
RP/0/RSP0/CPU0:PI# show running-config router bgp
Fri May 23 09:11:51.949 UTC
router bgp 64500
  address-family ipv4 unicast
  !
  address-family ipv6 unicast
    network 2001:db8:10:0:1::1/128
  !
  neighbor 10.0.2.1
    remote-as 64500
    update-source Loopback0
    address-family ipv4 unicast
  !
  !
  neighbor 10.1.1.1
    remote-as 64500
    update-source Loopback0
    address-family ipv4 unicast
    route-reflector-client
  !
  !
  neighbor 10.2.1.1
    remote-as 64500
    update-source Loopback0
    address-family ipv4 unicast
    route-reflector-client
  !
  !
  neighbor 10.3.1.1
    remote-as 64500
    update-source Loopback0
    address-family ipv4 unicast
    route-reflector-client
  !
  !
  neighbor 10.4.1.1
    remote-as 64500
    update-source Loopback0
    address-family ipv4 unicast
    route-reflector-client
  !
<... output omitted ...>
```

- Step 2** On the PE router, examine the EBGP and IBGP sessions configuration.

The PE router (Cisco IOS XR Software) output should be similar to the following, taken from Pod 3:

```
RP/0/RSP0/CPU0:PE3# show running-config router bgp
Fri May 23 08:46:10.716 UTC
router bgp 64500
  address-family ipv4 unicast
    network 10.3.1.1/32
  !
  address-family ipv6 unicast
    network 2001:db8:10:3:1::1/128
  !
  neighbor 10.0.1.1
    remote-as 64500
    update-source Loopback0
    address-family ipv4 unicast
      next-hop-self
    !
  !
  neighbor 192.168.103.31
    remote-as 64503
    address-family ipv4 unicast
      route-policy pass in
      route-policy pass out
      default-originate
<... output omitted ...>
```

The PE router (Cisco IOS XE Software) output should be similar to the following, taken from Pod 4:

```
PE4# show running-config | section router bgp
router bgp 64500
  bgp log-neighbor-changes
  neighbor 10.0.1.1 remote-as 64500
  neighbor 10.0.1.1 update-source Loopback0
  neighbor 2001:DB8:10:0:1::1 remote-as 64500
  neighbor 2001:DB8:10:0:1::1 update-source Loopback0
  neighbor 2001:DB8:192:168:104::41 remote-as 64504
  neighbor 192.168.104.41 remote-as 64504
  !
  address-family ipv4
    network 10.4.1.1 mask 255.255.255.255
    neighbor 10.0.1.1 activate
    neighbor 10.0.1.1 next-hop-self
    no neighbor 2001:DB8:10:0:1::1 activate
    no neighbor 2001:DB8:192:168:104::41 activate
    neighbor 192.168.104.41 activate
    neighbor 192.168.104.41 default-originate
  exit-address-family
<... output omitted ...>
```

Step 3 Verify that the EBGp session is established between the PE and CE routers in your pod. In the “state/prefix received” column, you should see a number other than zero.

The PE router (Cisco IOS XR Software) output should be similar to the following, taken from Pod 3:

```
RP/0/RSP0/CPU0:PE3# show bgp summary
Fri May 23 08:57:07.095 UTC
BGP router identifier 10.3.1.1, local AS number 64500
BGP generic scan interval 60 secs
BGP table state: Active
Table ID: 0xe0000000 RD version: 10
BGP main routing table version 10
BGP scan interval 60 secs
```

BGP is operating in STANDALONE mode.

Process	RcvTblVer	bRIB/RIB	LabelVer	ImportVer	SendTblVer	StandbyVer
Speaker	10	10	10	10	10	10

Neighbor	Spk	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	St/PfxRcd
10.0.1.1	0	64500	1174	1173	10	0	0	19:28:26	3
192.168.103.31	0	64503	1289	1176	10	0	0	19:28:28	1

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

The PE router (Cisco IOS XE Software) output should be similar to the following, taken from Pod 4:

```
PE4# show ip bgp summary
BGP router identifier 10.4.1.1, local AS number 64500
BGP table version is 7, main routing table version 7
5 network entries using 1240 bytes of memory
5 path entries using 600 bytes of memory
5/4 BGP path/bestpath attribute entries using 1200 bytes of memory
1 BGP rrinfo entries using 40 bytes of memory
2 BGP AS-PATH entries using 48 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 3128 total bytes of memory
BGP activity 11/0 prefixes, 11/0 paths, scan interval 60 secs
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
10.0.1.1	4	64500	1189	1305	7	0	0	19:42:36	3
192.168.104.41	4	64504	1308	1306	7	0	0	19:42:03	1

- Step 4** Verify that the IBGP session is established between the PE router in your pod and the P1 router. In the “state/prefix received” column, you should see a number other than zero.

The PE router (Cisco IOS XR Software) output should be similar to the following, taken from Pod 3:

```
RP/0/RSP0/CPU0:PE3# show bgp summary
Tue Jul 29 09:00:50.389 UTC
BGP router identifier 10.3.1.1, local AS number 64500
BGP generic scan interval 60 secs
BGP table state: Active
Table ID: 0xe0000000 RD version: 7
BGP main routing table version 7
BGP scan interval 60 secs
```

BGP is operating in STANDALONE mode.

Process	RcvTblVer	bRIB/RIB	LabelVer	ImportVer	SendTblVer	StandbyVer
Speaker	7	7	7	7	7	7

Neighbor	Spk	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	St/PfxRcd
10.0.1.1	0	64500	51	48	7	0	0	00:43:31	3
192.168.103.31	0	64503	78	74	7	0	0	01:07:11	1

The PE router (Cisco IOS XE Software) output should be similar to the following, taken from Pod 4:

```
PE4# show bgp summary
BGP router identifier 10.4.1.1, local AS number 64500
BGP table version is 7, main routing table version 7
5 network entries using 1240 bytes of memory
5 path entries using 600 bytes of memory
5/4 BGP path/bestpath attribute entries using 1200 bytes of memory
1 BGP rrinfo entries using 40 bytes of memory
2 BGP AS-PATH entries using 48 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 3128 total bytes of memory
BGP activity 11/0 prefixes, 11/0 paths, scan interval 60 secs
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
10.0.1.1	4	64500	1191	1308	7	0	0	19:44:46	3
192.168.104.41	4	64504	1310	1308	7	0	0	19:44:13	1

Step 5 On the PE router in your pod, determine if any routes are in the BGP table. You should see at least the route from the other pod in the team.

The PE router (Cisco IOS XR Software) output should be similar to the following, taken from Pod 3:

```
RP/0/RSP0/CPU0:PE3# show bgp
Tue Jul 29 10:38:59.503 UTC
BGP router identifier 10.3.1.1, local AS number 64500
BGP generic scan interval 60 secs
BGP table state: Active
Table ID: 0xe0000000 RD version: 7
BGP main routing table version 7
BGP scan interval 60 secs

Status codes: s suppressed, d damped, h history, * valid, > best
               i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network        Next Hop           Metric LocPrf Weight Path
*>i10.0.1.1/32     10.0.1.1             0      100      0 i
*> 10.3.1.1/32     0.0.0.0              0             32768 i
*> 10.3.10.1/32    192.168.103.31       0             0 64503 i
*>i10.4.1.1/32     10.4.1.1             0      100      0 i
*>i10.4.10.1/32    10.4.1.1             0      100      0 64504 i

Processed 5 prefixes, 5 paths
```

The PE router (Cisco IOS XE Software) output should be similar to the following, taken from Pod 4:

```
PE4# show ip bgp
BGP table version is 8, local router ID is 10.4.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
               x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

   Network        Next Hop           Metric LocPrf Weight Path
0.0.0.0           0.0.0.0              0             0 i
r>i 10.0.1.1/32    10.0.1.1             0      100      0 i
r>i 10.3.1.1/32    10.3.1.1             0      100      0 i
*>i 10.3.10.1/32   10.3.1.1             0      100      0 64503 i
*> 10.4.1.1/32     0.0.0.0              0             32768 i
*> 10.4.10.1/32    192.168.104.41       0             0 64504 i
```

- Step 6** On the PE router in your pod, verify the originator and cluster-ID list BGP attributes that have been inserted when the other pod route was reflected by the PI router. Your instructor configured the PI router as a route reflector for the PEx and PEy routers in your team.

The PE router (Cisco IOS XR Software) output should be similar to the following, taken from Pod 3:

```

RP/0/RSP0/CPU0:PE3# show bgp 10.4.10.1/32
Fri May 23 09:05:03.513 UTC
BGP routing table entry for 10.4.10.1/32
Versions:
  Process          BRIB/RIB  SendTblVer
  Speaker          9         9
Last Modified: May 22 13:38:19.829 for 19:26:43
Paths: (1 available, best #1)
  Advertised to peers (in unique update groups):
    192.168.103.31
  Path #1: Received by speaker 0
  Advertised to peers (in unique update groups):
    192.168.103.31
64504
  10.4.1.1 (metric 20) from 10.0.1.1 (10.4.1.1)
    Origin IGP, metric 0, localpref 100, valid, internal, best, group-best, import-
candidate
    Received Path ID 0, Local Path ID 1, version 9
    Originator: 10.4.1.1, Cluster list: 10.0.1.1

```

The PE router (Cisco IOS XE Software) output should be similar to the following, taken from Pod 4:

```

PE4# show ip bgp 10.3.10.1/32
BGP routing table entry for 10.3.10.1/32, version 4
Paths: (1 available, best #1, table default)
  Advertised to update-groups:
    2
Refresh Epoch 1
64503
  10.3.1.1 (metric 20) from 10.0.1.1 (10.0.1.1)
    Origin IGP, metric 0, localpref 100, valid, internal, best
    Originator: 10.3.1.1, Cluster list: 10.0.1.1
    rx pathid: 0, tx pathid: 0x0

```

Verification

No verification is needed for this task.

Task 2: Configure a Route Reflector and Internal BGP Session

In this task, you will configure the P2 router as a redundant route reflector. You will also configure internal BGP between the PE router in your pod and the P2 router in the backbone. You will also enable the BGP next-hop-self feature on the PE router.

Activity Procedure

Complete these steps:

- Step 1** IS-IS is used as the IGP in AS 64500. On the PE router in your pod, enable the interface that connects to the P2 router and enable IS-IS level-2 routing. IS-IS was configured by your instructor and is enabled on loopback interfaces and on links between the CE, PE, and P1 routers. Enable the interface and routing for both IPv4 and IPv6.
- Step 2** On the PE router in your pod, configure the internal BGP neighbor. The internal BGP neighbor is the P2 router in the backbone with IP address 10.0.2.1. The PE router in your pod should source BGP packets from the Loopback0 IP address.
- 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 P2 IBGP neighbor (BGP next-hop-self).
- Step 4** Configure the P2 router as an internal BGP neighbor to the PE router in your pod. Configure the P2 router as a route reflector for the PE router in your pod. Make sure that the cluster ID on the P2 router is set to the same value as on the P1 router. Coordinate this step with other teams if necessary.
- Step 5** Verify the IBGP sessions on the PE router in your pod. You should see P2 as the IBGP neighbor. The PE router (Cisco IOS XR Software) output should be similar to the following, taken from Pod 3:

```
RP/0/RSP0/CPU0:PE3# show bgp summary
Fri May 23 11:21:35.851 UTC
BGP router identifier 10.3.1.1, local AS number 64500
BGP generic scan interval 60 secs
BGP table state: Active
Table ID: 0xe0000000 RD version: 10
BGP main routing table version 10
BGP scan interval 60 secs
```

BGP is operating in STANDALONE mode.

Process	RcvTblVer		bRIB/RIB		LabelVer		ImportVer		SendTblVer		StandbyVer
Speaker	10		10		10		10		10		10
Neighbor	Spk	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	St/PfxRcd		
10.0.1.1	0	64500	1319	1317	10	0	0	21:52:55	3		
10.0.2.1	0	64500	9	5	10	0	0	00:00:37	3		
192.168.103.31	0	64503	1447	1321	10	0	0	21:52:56	1		

The PE router (Cisco IOS XE Software) output should be similar to the following, taken from Pod 4:

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

```
BGP router identifier 10.4.1.1, local AS number 64500
BGP table version is 7, main routing table version 7
5 network entries using 1240 bytes of memory
7 path entries using 840 bytes of memory
5/4 BGP path/bestpath attribute entries using 1200 bytes of memory
1 BGP rrinfo entries using 40 bytes of memory
2 BGP AS-PATH entries using 48 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 3368 total bytes of memory
BGP activity 11/0 prefixes, 13/0 paths, scan interval 60 secs
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
10.0.1.1	4	64500	1310	1440	7	0	0	21:44:12	3
10.0.2.1	4	64500	8	6	7	0	0	00:01:22	3
192.168.104.41	4	64504	1442	1440	7	0	0	21:43:40	1

Step 6 Verify the BGP table on the PE router. You should see the other pod route that is accessible over two paths. Note that the next hop in both cases is the originating router, the PE router in the other pod.

The PE router (Cisco IOS XR Software) output should be similar to the following, taken from Pod 3:

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

```
Tue Jul 29 11:11:26.925 UTC
BGP router identifier 10.3.1.1, local AS number 64500
BGP generic scan interval 60 secs
BGP table state: Active
Table ID: 0xe0000000 RD version: 7
BGP main routing table version 7
BGP scan interval 60 secs
```

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

Network	Next Hop	Metric	LocPrf	Weight	Path
*>i10.0.1.1/32	10.0.1.1	0	100	0	i
* i	10.0.1.1	0	100	0	i
*> 10.3.1.1/32	0.0.0.0	0		32768	i
*> 10.3.10.1/32	192.168.103.31	0		0	64503 i
*>i10.4.1.1/32	10.4.1.1	0	100	0	i
* i	10.4.1.1	0	100	0	i
*>i10.4.10.1/32	10.4.1.1	0	100	0	64504 i
* i	10.4.1.1	0	100	0	64504 i

```
Processed 5 prefixes, 8 paths
```

The PE router (Cisco IOS XE Software) output should be similar to the following, taken from Pod 4:

```

PE4# show ip bgp
BGP table version is 8, local router ID is 10.4.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
               x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

```

	Network	Next Hop	Metric	LocPrf	Weight	Path
	0.0.0.0	0.0.0.0			0	i
r i	10.0.1.1/32	10.0.1.1	0	100	0	i
r>i	10.0.1.1	10.0.1.1	0	100	0	i
r i	10.3.1.1/32	10.3.1.1	0	100	0	i
r>i	10.3.1.1	10.3.1.1	0	100	0	i
* i	10.3.10.1/32	10.3.1.1	0	100	0	64503 i
*>i	10.3.1.1	10.3.1.1	0	100	0	64503 i
*>	10.4.1.1/32	0.0.0.0	0		32768	i
*>	10.4.10.1/32	192.168.104.41	0		0	64504 i

Recall that the route reflector does not change the next-hop IP address when a route is reflected.

Step 7 Examine the details about the other pod route. You should see that this route actually originated from P1 and P2, respectively.

The PE router (Cisco IOS XR Software) output should be similar to the following, taken from Pod 3:

```

RP/0/RSP0/CPU0:PE3# show bgp 10.4.10.1/32
Fri May 23 11:30:12.599 UTC
BGP routing table entry for 10.4.10.1/32
Versions:
  Process          bRIB/RIB    SendTblVer
  Speaker          9           9
Last Modified: May 22 13:38:19.829 for 21:51:52
Paths: (2 available, best #1)
  Advertised to peers (in unique update groups):
    192.168.103.31
  Path #1: Received by speaker 0
  Advertised to peers (in unique update groups):
    192.168.103.31
  64504
    10.4.1.1 (metric 20) from 10.0.1.1 (10.4.1.1)
      Origin IGP, metric 0, localpref 100, valid, internal, best, group-best, import-
candidate
      Received Path ID 0, Local Path ID 1, version 9
      Originator: 10.4.1.1, Cluster list: 10.0.1.1
  Path #2: Received by speaker 0
  Not advertised to any peer
  64504
    10.4.1.1 (metric 20) from 10.0.2.1 (10.4.1.1)
      Origin IGP, metric 0, localpref 100, valid, internal
      Received Path ID 0, Local Path ID 0, version 0
      Originator: 10.4.1.1, Cluster list: 10.0.1.1

```

The PE router (Cisco IOS XE Software) output should be similar to the following, taken from Pod 4:

```

PE4# show ip bgp 10.3.10.1/32
BGP routing table entry for 10.3.10.1/32, version 4
Paths: (2 available, best #2, table default)
  Advertised to update-groups:
    2
  Refresh Epoch 1
  64503
    10.3.1.1 (metric 20) from 10.0.2.1 (10.0.2.1)
      Origin IGP, metric 0, localpref 100, valid, internal
      Originator: 10.3.1.1, Cluster list: 10.0.1.1
      rx pathid: 0, tx pathid: 0
  Refresh Epoch 1
  64503
    10.3.1.1 (metric 20) from 10.0.1.1 (10.0.1.1)
      Origin IGP, metric 0, localpref 100, valid, internal, best
      Originator: 10.3.1.1, Cluster list: 10.0.1.1
      rx pathid: 0, tx pathid: 0x0

```

Step 8 Answer the following question:

Which BGP path selection criterion is being used to select the path to the other pod route?

Verification

No verification is needed for this task.

Task 3: Restrict Route Propagation to a Client

In this task, you will restrict route propagation to route reflector clients on the P2 router. You will allow the P2 router to advertise only other pod routes to route reflector clients.

Activity Procedure

Complete these steps:

- Step 1** On the P2 router, create a route policy that will allow only a route originating from the other pod AS (for example, for Pod x, allow only routes that originate in AS 6450y). Use the **FILTER_TO_CLIENT_PODX** or **FILTER_TO_CLIENT_PODY** name as the route policy name.
- Step 2** On the P2 router, apply the route policy to the IBGP session with the pod PE router in the outbound direction.
- Step 3** On the PE router, verify the BGP table. You should see the routes from pods from other teams that are available only over the P1 router (because the P2 router reflects only routes from the other pod).

The PE router (Cisco IOS XR Software) output should be similar to the following, taken from Pod 3:

```

RP/0/RSP0/CPU0:PE3# show bgp
Tue Jul 29 11:21:27.370 UTC
BGP router identifier 10.3.1.1, local AS number 64500
BGP generic scan interval 60 secs
BGP table state: Active
Table ID: 0xe0000000 RD version: 7
BGP main routing table version 7
BGP scan interval 60 secs

Status codes: s suppressed, d damped, h history, * valid, > best
               i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network        Next Hop           Metric LocPrf Weight Path
*>i10.0.1.1/32     10.0.1.1             0      100      0 i
*> 10.3.1.1/32     0.0.0.0              0           32768 i
*> 10.3.10.1/32    192.168.103.31       0           0 64503 i
*>i10.4.1.1/32     10.4.1.1             0      100      0 i
*>i10.4.10.1/32    10.4.1.1             0      100      0 64504 i
* i               10.4.1.1             0      100      0 64504 i

Processed 5 prefixes, 6 paths

```

Routers from the other pod in the same team should be available over two paths.

The PE router (Cisco IOS XE Software) output should be similar to the following, taken from Pod 4:

```

PE4# show ip bgp
BGP table version is 8, local router ID is 10.4.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
               x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

   Network        Next Hop           Metric LocPrf Weight Path
0.0.0.0           0.0.0.0              0           0 i
r>i 10.0.1.1/32    10.0.1.1             0      100      0 i
r>i 10.3.1.1/32    10.3.1.1             0      100      0 i
* i 10.3.10.1/32   10.3.1.1             0      100      0 64503 i
*>i               10.3.1.1             0      100      0 64503 i
*> 10.4.1.1/32     0.0.0.0              0           32768 i
*> 10.4.10.1/32    192.168.104.41       0           0 64504 i

```

Verification

No verification is needed for this task.

Hardware Lab 2: Implement BGP Security Options

Activity Objective

In this lab activity, you will first configure BGP neighbor authentication between the CE, PE, and P1 routers. You will also enable a BGP TTL security check between the same routers. You will also configure CoPP on the CE router. Finally, you will implement source-based RTBH filtering by using the P1 router as an RTBH triggering router.


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

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

- Implement BGP neighbor authentication using a password
- Implement BGP neighbor authentication using key chains
- Enable the BGP TTL security check
- Enable CoPP
- Enable RTBH filtering

Visual Objective

The figure illustrates what you will accomplish in this activity.



Command List

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

Cisco IOS/IOS XE Software Commands

Command	Description
<code>class class_map_name</code>	Specifies the name of the class whose policy you want to create or change.
<code>class-map class_map_name</code>	Creates a class map.
<code>configure terminal</code>	Enters configuration mode.
<code>control-plane</code>	Enters control plane virtual interface configuration mode.
<code>interface interface</code>	Enters interface configuration mode.
<code>ip access-list extended acl_name</code>	Creates an extended access control list.
<code>ip address ip_address subnet_mask</code>	Sets a primary or secondary IPv4 address for an interface and the subnet mask.

Command	Description
ip route <i>network subnet_mask outgoing_interface</i>	Creates a static route.
ip verify unicast source reachable-via rx	Enables strict uRPF on an interface.
match access-group name <i>acl_name</i>	Specifies ACL matching criteria inside a class map.
neighbor ip_address password <i>password</i>	Enables BGP neighbor authentication.
neighbor ip_address ttl-security hops <i>hops_number</i>	Enables a BGP TTL security check.
network <i>network mask mask</i>	Advertises the network through BGP.
permit <i>protocol source [operator] [port] destination [operator] [port]</i>	Creates a permit ACL entry.
ping <i>dest_ip_source source_interface</i>	Verifies connectivity between the source IP and destination IP.
police rate <i>pps pps conform-action transmit exceed-action drop</i>	Configures traffic policing.
policy-map <i>policy_map_name</i>	Creates a policy map.
router bgp <i>AS-number</i>	Creates a BGP process and enters process configuration mode.
service-policy input <i>policy_map_name</i>	Applies a policy map to an interface in the inbound direction.
show access-lists	Displays access-list information
show class-map	Displays class-map information.
show ip bgp [<i>prefix</i>]	Displays the BGP routing table.
show ip bgp neighbors <i>ip_address</i>	Displays BGP neighbor information.
show policy-map control-plane	Displays the configuration either of a class or of all classes for the policy map of a control plane.

Cisco IOS XR Software Commands

Command	Description
accept-lifetime <i>start_time end_time</i>	Specifies key accept validity.

Command	Description
address-family <i>[ipv4 ipv6] unicast</i>	Enters address family configuration mode.
commit	Commits changes to the running configuration.
configure terminal	Enters configuration mode.
cryptographic-algorithm <i>algorithm</i>	Specifies the cryptographic algorithm for a key.
interface <i>interface</i>	Enters interface configuration mode.
ipv4 verify unicast source reachable-via rx	Enables strict uRPF on an interface.
key chain <i>keychain_name</i>	Creates a key chain.
key <i>key_id</i>	Specifies a key ID.
keychain <i>keychain_name</i>	Enables BGP neighbor authentication using a key chain.
neighbor <i>IP-address</i>	Configures the BGP neighbor and enters BGP neighbor configuration mode.
network <i>prefix outgoing_interface [tag tag]</i>	Configures the static route.
password <i>password</i>	Enables BGP neighbor authentication by using a password.
ping <i>dest_IP source source_IP</i>	Verifies connectivity between the source IP and destination IP (IPv4 and IPv6).
router bgp <i>AS-number</i>	Creates a BGP process and enters the BGP process configuration mode.
router static	Enters static route configuration mode.
send-lifetime <i>start_time end_time</i>	Specifies a key accept validity.
show bgp <i>[prefix]</i>	Displays the BGP routing table.
show bgp summary	Displays BGP routing protocol characteristics, including BGP neighbor status.
show lpts flows	Displays information about LPTS flows.
ttl-security	Enables BGP TTL security check.

Task 1: Implement BGP Neighbor Authentication Using Passwords

In this task, you will implement BGP neighbor authentication on the EBGP session between the CE and PE routers.

Activity Procedure

Complete these steps:

- Step 1** Verify that the EBGP session is established between the PE and CE routers in your pod. In the “state/prefix received” column, you should see a number other than zero.

The CE router (Cisco IOS Software) output should be similar to the following, taken from Pod 3:

```
CE3# show ip bgp summary
BGP router identifier 10.3.10.1, local AS number 64503
BGP table version is 10, main routing table version 10
5 network entries using 720 bytes of memory
5 path entries using 400 bytes of memory
4/4 BGP path/bestpath attribute entries using 576 bytes of memory
2 BGP AS-PATH entries using 48 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 1744 total bytes of memory
BGP activity 15/4 prefixes, 15/4 paths, scan interval 60 secs
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
192.168.103.30	4	64500	11238	12356	10	0	0	1w0d	5

The PE router (Cisco IOS XR Software) output should be similar to the following, taken from Pod 3:

```
RP/0/RSP0/CPU0:PE3# show bgp summary
Fri May 30 08:43:32.093 UTC
BGP router identifier 10.3.1.1, local AS number 64500
BGP generic scan interval 60 secs
BGP table state: Active
Table ID: 0xe0000000 RD version: 10
BGP main routing table version 10
BGP scan interval 60 secs
```

BGP is operating in STANDALONE mode.

Process	RcvTblVer	bRIB/RIB	LabelVer	ImportVer	SendTblVer	StandbyVer
Speaker	10	10	10	10	10	10

Neighbor	Spk	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	St/PfxRcd
10.0.1.1	0	64500	11240	11239	10	0	0	1w0d	3
10.0.2.1	0	64500	9941	9934	10	0	0	6d20h	1
192.168.103.31	0	64503	12361	11242	10	0	0	1w0d	1

The PE router (Cisco IOS XE Software) output should be similar to the following, taken from Pod 4:

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

```
BGP router identifier 10.4.1.1, local AS number 64500
BGP table version is 7, main routing table version 7
5 network entries using 1240 bytes of memory
6 path entries using 720 bytes of memory
5/4 BGP path/bestpath attribute entries using 1200 bytes of memory
1 BGP rrinfo entries using 40 bytes of memory
2 BGP AS-PATH entries using 48 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 3248 total bytes of memory
BGP activity 11/0 prefixes, 15/3 paths, scan interval 60 secs
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
10.0.1.1	4	64500	11232	12353	7	0	0	1w0d	3
10.0.2.1	4	64500	9890	10882	7	0	0	6d20h	1
192.168.104.41	4	64504	12361	12365	7	0	0	1w0d	1

Step 2 Enable BGP neighbor authentication on the CE router on the EBGP session with the PE router. Use `C!sc()` as a password.

On the CE router clear the BGP session because it will not be torn down automatically.

```
CE3# clear ip bgp *
```

```
CE3#
*May 30 08:50:37.125: %BGP-5-ADJCHANGE: neighbor 2001:DB8:192:168:103::30 Down User
reset
*May 30 08:50:37.125: %BGP_SESSION-5-ADJCHANGE: neighbor 2001:DB8:192:168:103::30 IPv6
Unicast topology base removed from session User reset
*May 30 08:50:37.125: %BGP-5-ADJCHANGE: neighbor 192.168.103.30 Down User reset
*May 30 08:50:37.125: %BGP_SESSION-5-ADJCHANGE: neighbor 192.168.103.30 IPv4 Unicast
topology base removed from session User reset
*May 30 08:50:37.753: %BGP-5-NBR RESET: Neighbor 2001:DB8:192
CE3#::168:103::30 active reset (Peer closed the session)
*May 30 08:50:37.753: %BGP_SESSION-5-ADJCHANGE: neighbor 2001:DB8:192:168:103::30 IPv6
Unicast topology base removed from session Peer closed the session
CE3#
```

Watch the CE router console. You should see that the CE router does not receive the MD5 hash from the PE router. Output should be similar to the following, taken from Pod 3:

```
CE3#
*May 30 08:50:52.661: %BGP-5-ADJCHANGE: neighbor 2001:DB8:192:168:103::30 Up
CE3#
*May 30 08:51:14.577: %TCP-6-BADAUTH: No MD5 digest from 192.168.103.30(18093) to
192.168.103.31(179) tableid = 0
```

Step 3 Enable BGP neighbor authentication on the PE router.

Verification

You have completed this task when you attain these results:

Step 1 Verify that an EBGP session is established between the PE and CE routers in your pod. In the “state/prefix received” column, you should see a number other than zero.

The CE router (Cisco IOS Software) output should be similar to the following, taken from Pod 3:

CE3# show ip bgp summary

BGP router identifier 10.3.10.1, local AS number 64503
BGP table version is 6, main routing table version 6
5 network entries using 720 bytes of memory
5 path entries using 400 bytes of memory
4/4 BGP path/bestpath attribute entries using 576 bytes of memory
2 BGP AS-PATH entries using 48 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 1744 total bytes of memory
BGP activity 27/16 prefixes, 27/16 paths, scan interval 60 secs

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
192.168.103.30	4	64500	24	25	6	0	0	00:18:24	5

The CE router (Cisco IOS Software) output should be similar to the following, taken from Pod 4:

CE4# show ip bgp summary

BGP router identifier 10.4.10.1, local AS number 64504
BGP table version is 7, main routing table version 7
6 network entries using 864 bytes of memory
6 path entries using 480 bytes of memory
4/4 BGP path/bestpath attribute entries using 576 bytes of memory
2 BGP AS-PATH entries using 48 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 1968 total bytes of memory
BGP activity 29/17 prefixes, 29/17 paths, scan interval 60 secs

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
192.168.104.40	4	64500	10	7	7	0	0	00:02:26	5

The PE router (Cisco IOS XR Software) output should be similar to the following, taken from Pod 3:

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

Fri May 30 09:26:37.646 UTC
BGP router identifier 10.3.1.1, local AS number 64500
BGP generic scan interval 60 secs
BGP table state: Active
Table ID: 0xe0000000 RD version: 15
BGP main routing table version 15
BGP scan interval 60 secs

BGP is operating in STANDALONE mode.

Process	RcvTblVer	bRIB/RIB	LabelVer	ImportVer	SendTblVer	StandbyVer
Speaker	15	15	15	15	15	15

Neighbor	Spk	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	St/PfxRcd
10.0.1.1	0	64500	11286	11284	15	0	0	1w0d	3
10.0.2.1	0	64500	9986	9979	15	0	0	6d21h	1
192.168.103.31	0	64503	12404	11284	15	0	0	00:19:00	1

The PE router (Cisco IOS XE Software) output should be similar to the following, taken from Pod 4:

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

```
BGP router identifier 10.4.1.1, local AS number 64500
BGP table version is 12, main routing table version 12
5 network entries using 1240 bytes of memory
6 path entries using 720 bytes of memory
5/4 BGP path/bestpath attribute entries using 1200 bytes of memory
1 BGP rrinfo entries using 40 bytes of memory
2 BGP AS-PATH entries using 48 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 3248 total bytes of memory
BGP activity 13/2 prefixes, 20/8 paths, scan interval 60 secs
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
10.0.1.1	4	64500	11279	12402	12	0	0	1w0d	3
10.0.2.1	4	64500	9936	10931	12	0	0	6d21h	1
192.168.104.41	4	64504	27	30	12	0	0	00:20:42	1

Task 2: Implement BGP Neighbor Authentication Using Key Chains

In this task, you will implement BGP neighbor authentication on the IBGP session between PE and P1 routers using key chains. BGP neighbor authentication using key chains is also called advanced BGP authentication and is not compatible with classic BGP authentication using passwords. This task applies to pods that are running the Cisco IOS XR Software router as the PE router only.

Activity Procedure

Complete these steps:

- Step 1** Verify that the IBGP session is established between the PE and P1 routers in your pod. In the “state/prefix received” column, you should see a number other than zero.

The PE router (Cisco IOS XR Software) output should be similar to the following, taken from Pod 3:

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

```
Fri May 30 10:07:49.205 UTC
BGP router identifier 10.3.1.1, local AS number 64500
BGP generic scan interval 60 secs
BGP table state: Active
Table ID: 0xe0000000 RD version: 15
BGP main routing table version 15
BGP scan interval 60 secs
```

BGP is operating in STANDALONE mode.

Process	RcvTblVer	bRIB/RIB	LabelVer	ImportVer	SendTblVer	StandbyVer
Speaker	15	15	15	15	15	15

Neighbor	Spk	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	St/PfxRcd
10.0.1.1	0	64500	11327	11325	15	0	0	1w0d	3
10.0.2.1	0	64500	10027	10020	15	0	0	6d22h	1
192.168.103.31	0	64503	12449	11325	15	0	0	01:00:11	1

- Step 2** Configure a key chain on the PE router. Use **C!sc()** as a key string and **HMAC-MD5** as a hashing algorithm. Also specify a valid send and accept lifetime; otherwise, the key will not be valid.
- Step 3** Apply the key chain on the PE router to the IBGP session with the P1 router. You should see that the PE router received packets with invalid authentication:

```
RP/0/RSP0/CPU0:PE3#
RP/0/RSP0/CPU0:May 30 10:18:28.171 : tcp[425]: %IP-TCP-3-BADAUTH : Invalid EA digest
from 10.0.1.1:54844 to 10.3.1.1:179
```

- Step 4** Use Telnet to connect to the P1 router. Configure a key chain on the P1 router. Use a pod-specific name for the key chain (for example, BGP_PODX or BGP_PODY). Use **C!sc()** as a key string and **HMAC-MD5** as a hashing algorithm. Also specify a valid send and accept lifetime; otherwise, the key will not be valid.
- Step 5** Apply the key chain on the P1 router to the IBGP session with the PE router in your pod. You should not see the message about the received packets with invalid authentication anymore.

Verification

You have completed this task when you attain these results:

- Step 1** Verify that an IBGP session is established between the PE and P1 routers in your pod. In the “state/prefix received” column, you should see a number other than zero.

The P1 router output should be similar to the following:

```
RP/0/RSP0/CPU0:P1# show bgp summary
Tue Jul 29 11:50:11.294 UTC
BGP router identifier 10.0.1.1, local AS number 64500
BGP generic scan interval 60 secs
BGP table state: Active
Table ID: 0xe0000000 RD version: 93
BGP main routing table version 93
BGP scan interval 60 secs

BGP is operating in STANDALONE mode.
```

Process Speaker	RcvTblVer 93	bRIB/RIB 93	LabelVer 93	ImportVer 93	SendTblVer 93	StandbyVer 93			
Neighbor	Spk	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	St/PfxRcd
10.0.2.1	0	64500	9527	9532	93	0	0	00:34:32	0
10.1.1.1	0	64500	0	0	0	0	0	00:00:00	Active
10.2.1.1	0	64500	0	0	0	0	0	00:00:00	Active
10.3.1.1	0	64500	3164	3189	93	0	0	00:00:10	2
10.4.1.1	0	64500	3479	3196	93	0	0	03:53:29	2

The PE router (Cisco IOS XR Software) output should be similar to the following, taken from Pod 3:

```
RP/0/RSP0/CPU0:PE3# show bgp summary
Fri May 30 10:29:06.298 UTC
BGP router identifier 10.3.1.1, local AS number 64500
BGP generic scan interval 60 secs
BGP table state: Active
Table ID: 0xe0000000 RD version: 19
BGP main routing table version 19
BGP scan interval 60 secs
```

BGP is operating in STANDALONE mode.

Process	RcvTblVer	bRIB/RIB	LabelVer	ImportVer	SendTblVer	StandbyVer
Speaker	19	19	19	19	19	19

Neighbor	Spk	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	St/PfxRcd
10.0.1.1	0	64500	11347	11346	19	0	0	00:03:19	3
10.0.2.1	0	64500	10048	10041	19	0	0	6d22h	1
192.168.103.31	0	64503	12473	11346	19	0	0	01:21:29	1

Task 3: Enable BGP TTL Security Check

In this task, you will enable a BGP TTL security check on the EBGP session between the CE and PE routers.

Activity Procedure

Complete these steps:

- Step 1** Enable TTL security check for the EBGP session on the CE router in your pod.
 - Step 2** What is the number that you have to specify with the command to enforce that EBGP neighbors are directly connected?
-
- Step 3** Wait 3 minutes. You should see that the EBGP session was torn down by the CE router:
The CE router output should be similar to the following, taken from Pod 3:

```
CE3#
*May 30 10:30:59.477: %BGP-3-NOTIFICATION: sent to neighbor 192.168.103.30 4/0 (hold
time expired) 0 bytes
*May 30 10:30:59.477: %BGP-5-NBR_RESET: Neighbor 192.168.103.30 reset (BGP Notification
sent)
*May 30 10:30:59.477: %BGP-5-ADJCHANGE: neighbor 192.168.103.30 Down BGP Notification
sent
*May 30 10:30:59.477: %BGP-SESSION-5-ADJCHANGE: neighbor 192.168.103.30 IPv4 Unicast
topology base removed from session BGP Notification sent
```

The CE router output should be similar to the following, taken from Pod 4:

```

CE4#
*May 30 10:25:43.277: %BGP-3-NOTIFICATION: sent to neighbor 192.168.104.40 4/0 (hold
time expired) 0 bytes
*May 30 10:25:43.277: %BGP-5-NBR_RESET: Neighbor 192.168.104.40 reset (BGP Notification
sent)
*May 30 10:25:43.277: %BGP-5-ADJCHANGE: neighbor 192.168.104.40 Down BGP Notification
sent
*May 30 10:25:43.277: %BGP_SESSION-5-ADJCHANGE: neighbor 192.168.104.40 IPv4 Unicast
topology base removed from session BGP Notification sent

```

Step 4 On the PE router, enable the TTL security check. You should see that the EBGP session went up:

The CE router (Cisco IOS Software) output should be similar to the following, taken from Pod 3:

```

CE3#
*May 30 10:38:32.236: %BGP-5-ADJCHANGE: neighbor 192.168.103.30 Up

```

The CE router (Cisco IOS Software) output should be similar to the following, taken from Pod 4:

```

CE4#
*May 30 10:33:01.605: %BGP-5-ADJCHANGE: neighbor 192.168.104.40 Up

```

Verification

You have completed this task when you attain these results:

Step 1 On the CE router, verify incoming and outgoing TTL settings for the neighbor PE.

The CE router output should be similar to the following, taken from Pod 3:

```

CE3# show ip bgp neighbors 192.168.103.30 | include TTL
Connection is ECN Disabled, Minimum incoming TTL 254, Outgoing TTL 255

```

The CE router output should be similar to the following, taken from Pod 4:

```

CE4# show ip bgp neighbors 192.168.104.40 | include TTL
Connection is ECN Disabled, Minimum incoming TTL 254, Outgoing TTL 255

```

Step 2 On the PE router, verify the expected TTL value for the BGP packets from the CE neighbor. On the PE router, verify the expected TTL value for the BGP packets from the CE neighbor.

On the PE (Cisco IOS XR Software) router, the output should be similar to the following, taken from Pod 3:

```
RP/0/RSP0/CPU0:PE3# show lpts flows | begin BGP
< output omitted >
```

```
-----
L3-proto      : IPV4(2)
L4-proto      : TCP(6)
VRF-ID        : default (0x60000000)
Local-IP      : any
Remote-IP     : 192.168.103.31
Local-Port    : 179
Remote-Port   : any
Interface     : any (0x0)
Flow-type     : BGP-cfg-peer
Min-TTL       : 255
Slice         : BGP4 FM
Flags         : 0x8 (in Pre-IFIB)
Location      : 0/RSP0/CPU0
Element References
location / count / scope
0/RSP0/CPU0 / 1 / LR
-----
```

On the PE (Cisco IOS XE Software) router, the output should be similar to the following, taken from Pod 4:

```
PE4# show ip bgp neighbors 192.168.104.41 | include TTL
Connection is ECN Disabled, Minimum incoming TTL 254, Outgoing TTL 255
```

Task 4: Enable CoPP

In this task, you will enable CoPP to rate-limit BGP traffic to the CE router.

Activity Procedure

Complete these steps:

- Step 1** On the CE router, configure a named ACL that will permit BGP traffic from the PE to the CE router. Verify the configured ACL.
- Step 2** On the CE router, create a class map that will refer to the previously configured ACL. Verify the configured class map.
- Step 3** On the CE router, create a policy map that will rate-limit BGP traffic from the PE to the CE router to 200 packets per second.
- Step 4** On the CE router, apply the configured policy map to the control plane virtual interface using the service policy. Verify the applied policy map.

Verification

You have completed this task when you attain these results:

- Step 1** On the CE router verify the configured ACL.
The output should be similar to the following, taken from Pod 3:

```
CE3# show access-lists
Extended IP access list BGP TRAFFIC
 10 permit tcp host 192.168.103.30 host 192.168.103.31 eq bgp
 20 permit tcp host 192.168.103.30 eq bgp host 192.168.103.31 (32 matches)
```

The output should be similar to the following, taken from Pod 4:

```
CE4# show access-lists
Extended IP access list BGP TRAFFIC
 10 permit tcp host 192.168.104.40 host 192.168.104.41 eq bgp
 20 permit tcp host 192.168.104.40 eq bgp host 192.168.104.41 (34 matches)
```

Step 2 On the CE router verify the configured class map.

The output should be similar to the following, taken from Pod 3:

```
CE3# show class-map
Class Map match-any class-default (id 0)
  Match any

Class Map match-all BGP_CLASS (id 1)
  Match access-group name BGP_TRAFFIC
```

The output should be similar to the following, taken from Pod 4:

```
CE4# show class-map
Class Map match-any class-default (id 0)
  Match any

Class Map match-all BGP_CLASS (id 1)
  Match access-group name BGP_TRAFFIC
```

Step 3 On the CE router verify the applied policy map.

The output should be similar to the following, taken from Pod 3:

```

CE3# show policy-map control-plane
Control Plane

Service-policy input: COPP

Class-map: BGP_CLASS (match-all)
  41 packets, 3578 bytes
  5 minute offered rate 0000 bps, drop rate 0000 bps
  Match: access-group name BGP_TRAFFIC
  police:
    rate 200 pps, burst 48 packets
    conformed 41 packets, 41 bytes; actions:
      transmit
    exceeded 0 packets, 0 bytes; actions:
      drop
    conformed 0 pps, exceeded 0 pps

Class-map: class-default (match-any)
  106 packets, 21069 bytes
  5 minute offered rate 0000 bps, drop rate 0000 bps
  Match: any

```

The output should be similar to the following, taken from Pod 4:

```

CE4# show policy-map control-plane
Control Plane

Service-policy input: COPP

Class-map: BGP_CLASS (match-all)
  43 packets, 3600 bytes
  5 minute offered rate 0000 bps, drop rate 0000 bps
  Match: access-group name BGP_TRAFFIC
  police:
    rate 200 pps, burst 48 packets
    conformed 43 packets, 43 bytes; actions:
      transmit
    exceeded 0 packets, 0 bytes; actions:
      drop
    conformed 0 pps, exceeded 0 pps

Class-map: class-default (match-any)
  71 packets, 12140 bytes
  5 minute offered rate 0000 bps, drop rate 0000 bps
  Match: any

```

Task 5: Enable RTBH Filtering

In this optional task, you will enable source-based RTBH filtering.

Activity Procedure

Complete these steps:

- Step 1** On the CE router, create the Loopback1 interface. Use 10.x.100.1/32 or 10.y.100.1/32 as the IP address on the interface. This interface will be used as a source of traffic that you want to black-hole.
- Step 2** On the CE router, advertise the previously configured /32 network on the loopback interface into BGP.
- Step 3** From the CE router, ping the other pod CE router from the Loopback1 interface. You should be successful.

The CE router output should be similar to the following, taken from Pod 3:

```
CE3# ping 10.4.100.1 source Loopback1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.4.100.1, timeout is 2 seconds:
Packet sent with a source address of 10.3.100.1
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

The CE router output should be similar to the following, taken from Pod 4:

```
CE4# ping 10.3.100.1 source Loopback1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.3.100.1, timeout is 2 seconds:
Packet sent with a source address of 10.4.100.1
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

- Step 4** On the PE router, create a static route for the 172.16.x.0/24 (or 172.16.y.0/24) network that points to the null0 interface.
- Step 5** On the PE router, enable strict uRPF on the CE-facing interface. Use the Job Aids section to determine the interface.
- Step 6** Use Telnet to connect to the P1 router. Create a static route for the 172.16.x.0/24 (or 172.16.y.0/24) network that points to the null0 interface.
- Step 7** On the P1 router, examine the preconfigured route policy named RTBH:

```
RP/0/RSP0/CPU0:P1# show running-config route-policy RTBH
Fri May 30 12:02:25.529 UTC
route-policy RTBH
  if tag eq 1 then
    set next-hop 172.16.1.1
    set local-preference 1000
    set community (no-export)
  elseif tag eq 2 then
    set next-hop 172.16.2.1
    set local-preference 1000
    set community (no-export)
  elseif tag eq 3 then
    set next-hop 172.16.3.1
    set local-preference 1000
    set community (no-export)
  elseif tag eq 4 then
    set next-hop 172.16.4.1
    set local-preference 1000
    set community (no-export)
  else
    drop
  endif
end-policy
```

- Step 8** On the P1 router, examine the BGP configuration. You should see that static routes are redistributed into BGP using the RTBH route policy as a filter:

```
RP/0/RSP0/CPU0:P1# show running-config router bgp
Fri May 30 12:05:48.727 UTC
router bgp 64500
  address-family ipv4 unicast
    redistribute static route-policy RTBH
<... output omitted ...>
```

- Step 9** Why are the redistributed routes tagged with no-export community?
-
- Step 10** Why is the local preference of redistributed routes set to 1000?
-
- Step 11** On the P1 router, trigger black-holing of traffic originating from the CE Loopback1 interface. Configure a static route for the Loopback1 interface that is tagged with x (or y) and points to the null0 interface.

Verification

You have completed this task when you attain these results:

- Step 1** On the PE router, examine the BGP table. You should see an additional path to the 10.x.100.1/32 (or 10.y.100.1/32) network that points to the 172.16.x.1 (or 172.16.y.1) next hop.

The PE (Cisco IOS XR) router output should be similar to the following, taken from Pod 3:

```
RP/0/RSP0/CPU0:PE3# show bgp
Tue Jul 29 13:53:22.772 UTC
BGP router identifier 10.3.1.1, local AS number 64500
BGP generic scan interval 60 secs
BGP table state: Active
Table ID: 0xe0000000 RD version: 28
BGP main routing table version 28
BGP scan interval 60 secs

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

   Network          Next Hop          Metric LocPrf Weight Path
*>i10.0.1.1/32      10.0.1.1              0      100      0 i
*> 10.3.1.1/32      0.0.0.0                0                 32768 i
*> 10.3.10.1/32     192.168.103.31         0                 0 64503 i
*> 10.3.100.1/32    192.168.103.31         0                 0 64503 i
*>i10.4.1.1/32      10.4.1.1              0      100      0 i
*>i10.4.10.1/32     10.4.1.1              0      100      0 64504 i
* i                10.4.1.1              0      100      0 64504 i
*>i10.4.100.1/32    10.4.1.1              0      100      0 64504 i
* i                10.4.1.1              0      100      0 64504 i

Processed 7 prefixes, 9 paths
```

The PE (Cisco IOS XE) router output should be similar to the following, taken from Pod 4:

```
PE4# show ip bgp
BGP table version is 28, local router ID is 10.4.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

   Network          Next Hop          Metric LocPrf Weight Path
0.0.0.0             0.0.0.0                0                 0 i
r>i 10.0.1.1/32      10.0.1.1              0      100      0 i
r>i 10.3.1.1/32      10.3.1.1              0      100      0 i
*>i 10.3.10.1/32     10.3.1.1              0      100      0 64503 i
* i                10.3.1.1              0      100      0 64503 i
*>i 10.3.100.1/32    10.3.1.1              0      100      0 64503 i
* i                10.3.1.1              0      100      0 64503 i
*> 10.4.1.1/32       0.0.0.0                0                 32768 i
*> 10.4.10.1/32      192.168.104.41         0                 0 64504 i
*> 10.4.100.1/32     192.168.104.41         0                 0 64504 i
```

- Step 2** On the PE router, examine the detailed Cisco Express Forwarding entry for the 10.x.100.1/32 (or 10.y.100.1/32) network.

The PE (Cisco IOS XR) router output should be similar to the following, taken from Pod 3:

```
RP/0/RSP0/CPU0:PE3# show cef 10.3.100.1/32 detail
Fri May 30 12:21:37.926 UTC
10.3.100.1/32, version 139, internal 0x14000001 0x0 {ptr 0xadd56f90} [1], 0x0 {0x0},
0x0 {0x0}
Updated May 30 11:40:07.850
Prefix Len 32, traffic index 0, precedence n/a, priority 4
gateway array {0xad32b7c0} reference count 2, flags 0x8020, source rib {6}, 0 backups
[1 type 3 flags 0x90111 {0xad488aa0} ext 0x0 {0x0}]
LW-LDI[type=0, refc=0, ptr=0x0, sh-ldi=0x0]

Level 1 - Load distribution: 0
[0] via 192.168.103.31, recursive

via 192.168.103.31, 3 dependencies, recursive, bgp-ext {flags 0x6020}
path-idx 0 NHID 0x0 {0xadd5693c 0x0}
next hop 192.168.103.31 via 192.168.103.31/32

Load distribution: 0 {refcount 1}

Hash OK Interface Address
0 Y GigabitEthernet0/0/0/0 remote
```

The PE (Cisco IOS XE) router output should be similar to the following, taken from Pod 4:

```
PE4# show ip cef 10.4.100.1/32 detail
10.4.100.1/32, epoch 2, flags rib only nolabel, rib defined all labels
recursive via 192.168.104.41
attached to GigabitEthernet0/0/0
```

Step 3 From the CE router, ping the other pod CE router. Use the Loopback1 interface as the source interface.

The CE router output should be similar to the following, taken from Pod 3:

```
CE3# ping 10.4.10.1 source Loopback1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.4.10.1, timeout is 2 seconds:
Packet sent with a source address of 10.3.100.1
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

The CE router output should be similar to the following, taken from Pod 4:

```
CE4# ping 10.3.10.1 source Loopback1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.3.10.1, timeout is 2 seconds:
Packet sent with a source address of 10.4.100.1
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

Hardware Lab 3: Improve BGP Scalability

Activity Objective

In this activity, you will first migrate an existing BGP configuration for EBGP neighbors to a template-based configuration. Then you will limit the number of prefixes that can be received from a BGP neighbor. You will also improve BGP convergence by changing the BGP scan and advertisement interval. Finally, you will implement BGP route dampening. All configurations will be performed on the PE router.


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

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

- Implement BGP configuration and peer templates
- Limit the number of prefixes that are received from a BGP neighbor
- Improve BGP convergence by changing the BGP scan and advertisement interval
- Implement BGP route dampening

Visual Objective

The figure illustrates what you will accomplish in this activity.



Command List

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

Cisco IOS/IOS XE Software Commands

Command	Description
[no] shutdown	Enables and disables an interface.
address-family [ipv4 ipv6]	Enters address family configuration mode inside the BGP process.
bfd interval <i>send_timer</i> min_rx <i>receive_timer</i> multiplier <i>number</i>	Enables BFD on an interface.
bgp dampening	Enables BGP dampening with default parameters.
bgp scan-time <i>scan_time</i>	Configures scanning intervals of BGP routers for next-hop validation.
configure terminal	Enters configuration mode.
debug ip bgp dampening	Enables BGP dampening debugging.

Command	Description
interface <i>interface</i>	Enters interface configuration mode.
maximum-prefix <i>num_of_prefixes</i>	Enables the maximum prefix feature inside a peer policy template.
neighbor ip_address advertisement-interval <i>advertisement_interval</i>	Changes the advertisement interval for a neighbor.
neighbor ip_address fall-over bfd	Enables BFD support for BGP.
neighbor ip_address inherit peer-policy <i>peer_policy_name</i>	Sends a peer policy template to a neighbor so that the neighbor can inherit the configuration.
neighbor ip_address inherit peer-session <i>peer_session_name</i>	Sends a peer session template to a neighbor so that the neighbor can inherit the configuration.
neighbor ip_address remote-as <i>remote_as</i>	Adds an entry to the BGP or MP-BGP neighbor table.
password <i>password</i>	Enables BGP authentication inside a peer session template.
router bgp <i>AS-number</i>	Creates a BGP process and enters BGP process configuration mode.
show ip bgp [<i>prefix</i>]	Displays the BGP routing table.
show ip bgp dampening dampened-paths	Displays BGP dampened routes.
show ip bgp neighbors <i>ip_address</i> [<i>configuration</i>]	Displays BGP neighbor information.
show ip bgp summary	Displays the status of all BGP connections.
show ip bgp template peer-policy	Displays locally configured peer policy templates.
show ip bgp template peer-session	Displays locally configured peer session templates.
switchport access vlan <i>vlan_id</i>	Changes the VLAN for a switch port.
template peer-policy <i>peer_policy_name</i>	Creates a peer policy template and enters policy-template configuration mode.
template peer-session <i>peer_session_name</i>	Creates a peer session template and enters session-template configuration mode.
ttl-security hops <i>hops</i>	Enables TTL security inside a peer session template.
undebg all	Disables all debugging.

Cisco IOS XR Software Commands

Command	Description
address-family [ipv4 ipv6] unicast	Enters address family configuration mode.
af-group <i>name</i> address-family ipv4 unicast	Creates an address family group for BGP neighbors and enters address family group configuration mode.
bfd fast-detect	Enables BFD support under router BGP configuration mode.
bfd minimum-interval <i>interval</i>	Sets the BFD minimum interval value under router BGP configuration mode.
bfd multiplier <i>number</i>	Sets the BFD multiplier value under router BGP configuration mode.
bgp dampening	Enables BGP dampening with default parameters.
bgp scan-time <i>scan_time</i>	Configures the scanning intervals of BGP routers for next-hop validation.
commit	Commits changes to the running configuration.
configure terminal	Enters configuration mode.
debug bgp dampening	Enables BGP dampening debugging.
maximum-prefix <i>num_of_prefixes</i>	Enables the maximum prefix feature inside a peer policy template.
neighbor <i>IP_address</i>	Configures the BGP neighbor and enters BGP neighbor configuration mode.
neighbor-group <i>name</i>	Creates a neighbor group and enters neighbor group configuration mode.
password <i>password</i>	Enables BGP neighbor authentication using a password.
router bgp <i>AS-number</i>	Creates a BGP process and enters BGP process configuration mode.
show bgp [<i>prefix</i>]	Displays the BGP routing table.
show bgp address-family [ipv4 ipv6] unicast dampened-paths	Displays BGP dampened routes.
show bgp af-group <i>af_group_name</i> configuration	Displays effective BGP configuration for address family groups.
show bgp neighbor-group <i>neighbor_group_name</i> configuration	Displays effective BGP configuration for neighbor groups.

Command	Description
<code>show bgp neighbors ip_address [configuration]</code>	Displays BGP neighbor information.
<code>show bgp summary</code>	Displays BGP routing protocol characteristics, including BGP neighbor status.
<code>ttl-security</code>	Enables BGP TTL security check.
<code>undebg all</code>	Disables all debugging.
<code>use af-group af_group_name</code>	Inherits a configuration from an address family group.
<code>use neighbor-group neighbor_group_name</code>	Inherits a configuration from a neighbor group.

Task 1: Implement the BGP Configuration and Peer Templates

In this task, you will implement the BGP configuration and peer templates on the PE router for the existing EBGP session with the CE router.

Activity Procedure

Complete these steps:

- Step 1** Verify that an EBGP session is established between the PE and CE routers in your pod. In the “state/prefix received” column, you should see a number other than zero.

The PE (Cisco IOS XR) router output should be similar to the following, taken from Pod 3:

```
RP/0/RSP0/CPU0:PE3# show bgp summary
Tue Jun 10 09:11:18.705 UTC
BGP router identifier 10.3.1.1, local AS number 64500
BGP generic scan interval 60 secs
BGP table state: Active
Table ID: 0xe0000000 RD version: 27
BGP main routing table version 27
BGP scan interval 60 secs

BGP is operating in STANDALONE mode.
```

Process Speaker	RcvTblVer 27	bRIB/RIB 27	LabelVer 27	ImportVer 27	SendTblVer 27	StandbyVer 27			
Neighbor	Spk	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	St/PfxRcd
10.0.1.1	0	64500	27113	27109	27	0	0	1w3d	4
10.0.2.1	0	64500	25812	25804	27	0	0	2w3d	2
192.168.103.31	0	64503	29815	27113	27	0	0	1w3d	2

The CE (Cisco IOS) router output should be similar to the following, taken from Pod 3:

CE3# show ip bgp summary

BGP router identifier 10.3.10.1, local AS number 64503
BGP table version is 16, main routing table version 16
7 network entries using 1008 bytes of memory
7 path entries using 560 bytes of memory
4/4 BGP path/bestpath attribute entries using 576 bytes of memory
2 BGP AS-PATH entries using 48 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 2192 total bytes of memory
BGP activity 33/20 prefixes, 33/20 paths, scan interval 60 secs

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
192.168.103.30	4	64500	15751	17327	16	0	0	1w3d	6

The PE (Cisco IOS XE) router output should be similar to the following, taken from Pod 4:

PE4# show ip bgp summary

BGP router identifier 10.4.1.1, local AS number 64500
BGP table version is 24, main routing table version 24
7 network entries using 1736 bytes of memory
9 path entries using 1080 bytes of memory
5/4 BGP path/bestpath attribute entries using 1200 bytes of memory
1 BGP rrinfo entries using 40 bytes of memory
2 BGP AS-PATH entries using 48 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 4104 total bytes of memory
BGP activity 18/5 prefixes, 28/13 paths, scan interval 60 secs

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
10.0.1.1	4	64500	27110	29817	24	0	0	2w4d	4
10.0.2.1	4	64500	25762	28346	24	0	0	2w3d	2
192.168.104.41	4	64504	17328	17326	24	0	0	1w3d	2

The CE (Cisco IOS) router output should be similar to the following, taken from Pod 4:

CE4# show ip bgp summary

BGP router identifier 10.4.10.1, local AS number 64504
BGP table version is 18, main routing table version 18
7 network entries using 1008 bytes of memory
7 path entries using 560 bytes of memory
4/4 BGP path/bestpath attribute entries using 576 bytes of memory
2 BGP AS-PATH entries using 48 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 2192 total bytes of memory
BGP activity 38/25 prefixes, 39/26 paths, scan interval 60 secs

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
192.168.104.40	4	64500	17326	17329	18	0	0	1w3d	6

Step 2 On the PE router, verify the BGP configuration for the CE neighbor.

The PE router output (Cisco IOS XR) should be similar to the following, taken from Pod 3:

```

RP/0/RSP0/CPU0:PE3# show running-config router bgp
Tue Jun 10 09:16:30.401 UTC
router bgp 64500
  address-family ipv4 unicast
    network 10.3.1.1/32
  !
  address-family ipv6 unicast
    network 2001:db8:10:3:1::1/128
  !
  neighbor 10.0.1.1
    remote-as 64500
    keychain BGP
    update-source Loopback0
    address-family ipv4 unicast
      next-hop-self
    !
  !
  neighbor 10.0.2.1
    remote-as 64500
    update-source Loopback0
    address-family ipv4 unicast
      next-hop-self
    !
  !
  neighbor 192.168.103.31
    remote-as 64503
    password encrypted 112A5816145A42
    ttl-security
    address-family ipv4 unicast
      route-policy pass in
      route-policy pass out
      default-originate
    !
  !
  neighbor 2001:db8:10:0:1::1
    remote-as 64500
    update-source Loopback0
    address-family ipv6 unicast
      next-hop-self
    !
  !
  neighbor 2001:db8:192:168:103::31
    remote-as 64503
    address-family ipv6 unicast
      route-policy pass in
      route-policy pass out
      default-originate

```

The PE router output (Cisco IOS XE) should be similar to the following, taken from Pod 4:

```

PE4# show running-config | section router bgp
router bgp 64500
  bgp log-neighbor-changes
  neighbor 10.0.1.1 remote-as 64500
  neighbor 10.0.1.1 update-source Loopback0
  neighbor 10.0.2.1 remote-as 64500
  neighbor 10.0.2.1 update-source Loopback0
  neighbor 2001:DB8:10:0:1::1 remote-as 64500
  neighbor 2001:DB8:10:0:1::1 update-source Loopback0
  neighbor 2001:DB8:192:168:104::41 remote-as 64504
  neighbor 192.168.104.41 remote-as 64504
  neighbor 192.168.104.41 ttl-security hops 1
  neighbor 192.168.104.41 password C!sc()
  !
  address-family ipv4
    network 10.4.1.1 mask 255.255.255.255
    neighbor 10.0.1.1 activate
    neighbor 10.0.1.1 next-hop-self
    neighbor 10.0.2.1 activate
    neighbor 10.0.2.1 next-hop-self
    no neighbor 2001:DB8:10:0:1::1 activate
    no neighbor 2001:DB8:192:168:104::41 activate
    neighbor 192.168.104.41 activate
    neighbor 192.168.104.41 default-originate
  exit-address-family
  !
  address-family ipv6
    network 2001:DB8:10:4:1::1/128
    neighbor 2001:DB8:10:0:1::1 activate
    neighbor 2001:DB8:10:0:1::1 next-hop-self
    neighbor 2001:DB8:192:168:104::41 activate
    neighbor 2001:DB8:192:168:104::41 default-originate
  exit-address-family

```

Step 3 On the PE router that is running Cisco IOS XR Software only, complete these steps:

- Migrate all session-specific commands for the EBGp neighbor into the neighbor group (except the **remote-as** command). Use **EBGP** as the name of the neighbor group.
- Migrate all address family-specific commands for the EBGp neighbor into the address family group. Use **IPv4** as the name of the address family group.
- Configure the EBGp neighbor group to inherit the configuration from the IPV4 address family group.
- Delete the CE router as an EBGp neighbor. Add the CE router as an EBGp neighbor again and configure the neighbor to inherit the configuration from the neighbor group.

Step 4 On the PE router that is running Cisco IOS XE Software only, complete these steps:

- Migrate all session-specific commands for the EBGp neighbor into the peer session template (except the **remote-as** command). Use **EBGP_SESSION** as the name of the peer session template.
- Migrate all address family-specific commands for the EBGp neighbor into the peer policy template. If no address family-specific configuration is present, create an empty peer policy template because it will be used in the next tasks. Use **EBGP_POLICY** as the name of the peer policy template.
- Delete the CE router as an EBGp neighbor. Add the CE router as an EBGp neighbor again and configure the neighbor to inherit the configuration from the peer session and peer policy templates.

Verification

You have completed this task when you attain these results:

- Step 1** On the PE router that is running Cisco IOS XR Software, verify the configured address family group.

The PE router (Cisco IOS XR) output should be similar to the following, taken from Pod 3:

```
RP/0/RSP0/CPU0:PE3# show bgp af-group IPv4 configuration
af-group IPv4 address-family IPv4 Unicast
  policy pass in          []
  policy pass out         []
```

- Step 2** On the PE router that is running Cisco IOS XR Software, verify the configured neighbor group. You should see the inherited configuration from the address family group.

The PE router (Cisco IOS XR) output should be similar to the following, taken from Pod 3:

```
RP/0/RSP0/CPU0:PE3# show bgp neighbor-group EBGp configuration
neighbor-group EBGp
  password encrypted 143453180F4C63 []
  ttl-security        []
  address-family IPv4 Unicast      []
  policy pass in       [a:IPv4]
  policy pass out      [a:IPv4]
```

- Step 3** On the PE router that is running Cisco IOS XR Software, verify the CE neighbor configuration. You should see the inherited configuration from the neighbor group.

The PE router (Cisco IOS XR) output should be similar to the following, taken from Pod 3:

```
RP/0/RSP0/CPU0:PE3# show bgp neighbors 192.168.103.31 configuration
neighbor 192.168.103.31
  remote-as 64503          []
  password encrypted 143453180F4C63 [n:EBGP]
  ttl-security             [n:EBGP]
  address-family IPv4 Unicast [n:EBGP]
  policy pass in           [n:EBGP a:IPv4]
  policy pass out          [n:EBGP a:IPv4]
```

- Step 4** On the PE router that is running Cisco IOS XE Software, verify the configured peer policy template.

The PE router (Cisco IOS XE) output should be similar to the following, taken from Pod 4:

```
PE4# show ip bgp template peer-policy
Template:EBGP POLICY, index:1.
Local policies:0x0, Inherited polices:0x0
Local disable policies:0x0, Inherited disable policies:0x0
Locally configured policies:
Inherited policies:
```

Step 5 On the PE router that is running Cisco IOS XE Software, verify the configured peer session template.

The PE router (Cisco IOS XE) output should be similar to the following, taken from Pod 4:

```
PE4# show ip bgp template peer-session
Template:EBGP_SESSION, index:1
Local policies:0x810, Inherited policies:0x0
Locally configured session commands:
  password is configured
  ttl-security hops 1
Inherited session commands:
```

Step 6 Verify that the EBGP session is established between the PE and CE routers in your pod. In the "state/prefix received" column, you should see a number other than zero.

The PE router output (Cisco IOS XR) should be similar to the following, taken from Pod 3:

```
RP/0/RSP0/CPU0:PE3# show bgp summary
Tue Jun 10 11:12:43.660 UTC
BGP router identifier 10.3.1.1, local AS number 64500
BGP generic scan interval 60 secs
BGP table state: Active
Table ID: 0xe0000000 RD version: 62
BGP main routing table version 62
BGP scan interval 60 secs

BGP is operating in STANDALONE mode.
```

Process	RcvTblVer	bRIB/RIB	LabelVer	ImportVer	SendTblVer	StandbyVer
Speaker	62	62	62	62	62	62

Neighbor	Spk	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	St/PfxRcd
10.0.1.1	0	64500	27252	27246	62	0	0	00:18:53	4
10.0.2.1	0	64500	25946	25941	62	0	0	00:18:55	2
192.168.103.31	0	64503	29966	27254	62	0	0	00:19:23	2

The CE router output (Cisco IOS) should be similar to the following, taken from Pod 3:

```
CE3# show ip bgp summary
BGP router identifier 10.3.10.1, local AS number 64503
BGP table version is 33, main routing table version 33
6 network entries using 864 bytes of memory
6 path entries using 480 bytes of memory
4/4 BGP path/bestpath attribute entries using 576 bytes of memory
2 BGP AS-PATH entries using 48 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 1968 total bytes of memory
BGP activity 37/25 prefixes, 56/44 paths, scan interval 60 secs
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
192.168.103.30	4	64500	25	27	33	0	0	00:19:55	5

The PE router output (Cisco IOS XE) should be similar to the following, taken from Pod 4:

```

PE4# show ip bgp summary
BGP router identifier 10.4.1.1, local AS number 64500
BGP table version is 44, main routing table version 44
6 network entries using 1488 bytes of memory
8 path entries using 960 bytes of memory
4/4 BGP path/bestpath attribute entries using 960 bytes of memory
1 BGP rrinfo entries using 40 bytes of memory
2 BGP AS-PATH entries using 48 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 3496 total bytes of memory
BGP activity 30/18 prefixes, 47/33 paths, scan interval 60 secs

```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
10.0.1.1	4	64500	27236	29950	44	0	0	2w4d	4
10.0.2.1	4	64500	25883	28479	44	0	0	2w3d	2
192.168.104.41	4	64504	114	118	44	0	0	01:39:39	2

The CE router output (Cisco IOS) should be similar to the following, taken from Pod 4:

```

CE4# show ip bgp summary
BGP router identifier 10.4.10.1, local AS number 64504
BGP table version is 37, main routing table version 37
6 network entries using 864 bytes of memory
6 path entries using 480 bytes of memory
4/4 BGP path/bestpath attribute entries using 576 bytes of memory
2 BGP AS-PATH entries using 48 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 1968 total bytes of memory
BGP activity 53/41 prefixes, 54/42 paths, scan interval 60 secs

```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
192.168.104.40	4	64500	119	114	37	0	0	01:40:03	5

Task 2: Limit the Number of Prefixes Received from a BGP Neighbor

In this task, you will enable the maximum prefixes feature on the PE router. The configuration will be added to the template that was configured in the previous task.

Activity Procedure

Complete these steps:

Step 1 On the PE router, examine how many routes are received from the CE router.

The PE router (Cisco IOS XR) output should be similar to the following, taken from Pod 3:

```
RP/0/RSP0/CPU0:PE3# show bgp summary
BGP router identifier 10.3.1.1, local AS number 64500
BGP generic scan interval 60 secs
BGP table state: Active
Table ID: 0xe0000000 RD version: 62
BGP main routing table version 62
BGP scan interval 60 secs
```

BGP is operating in STANDALONE mode.

Process Speaker	RcvTblVer	bRIB/RIB	LabelVer	ImportVer	SendTblVer	StandbyVer
	62	62	62	62	62	62

Neighbor	Spk	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	St/PfxRcd
10.0.1.1	0	64500	27263	27257	62	0	0	00:30:00	4
10.0.2.1	0	64500	25957	25952	62	0	0	00:30:02	2
192.168.103.31	0	64503	29979	27265	62	0	0	00:30:30	2

The PE router (Cisco IOS XE) output should be similar to the following, taken from Pod 4:

```
PE4# show ip bgp summary
BGP router identifier 10.4.1.1, local AS number 64500
BGP table version is 44, main routing table version 44
6 network entries using 1488 bytes of memory
8 path entries using 960 bytes of memory
4/4 BGP path/bestpath attribute entries using 960 bytes of memory
1 BGP rrinfo entries using 40 bytes of memory
2 BGP AS-PATH entries using 48 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 3496 total bytes of memory
BGP activity 30/18 prefixes, 47/33 paths, scan interval 60 secs
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
10.0.1.1	4	64500	27248	29962	44	0	0	2w4d	4
10.0.2.1	4	64500	25895	28492	44	0	0	2w3d	2
192.168.104.41	4	64504	126	131	44	0	0	01:50:58	2

- Step 2** On the PE router, enable the maximum prefix feature for routes that are received from EBGp neighbors. Configure the feature in the previously configured address family group or in the peer policy template (depending on the software that is running on the PE router). The maximum number of allowed prefixes should be large enough to accommodate all of the routes that are received from the CE router.

Verification

You have completed this task when you attain these results:

- Step 1** Verify the configuration of the maximum prefix for the CE neighbor. The PE router output should be similar to the following, taken from Pod 1:

The PE router (Cisco IOS XR) output should be similar to the following, taken from Pod 3:

```
RP/0/RSP0/CPU0:PE3# show bgp neighbors 192.168.103.31 | include Maximum
Maximum prefixes allowed 2
```

The PE router (Cisco IOS XE) output should be similar to the following, taken from Pod 4:

```
PE4# show ip bgp neighbors 192.168.104.41 | include Maximum
Maximum prefixes allowed 2
Maximum output segment queue size: 50
```

Task 3: Improve BGP Convergence by Changing the BGP Scan and Advertisement Interval

In this task, you will improve BGP convergence by changing the BGP scan and advertisement interval on the PE router.

Activity Procedure

Complete these steps:

Step 1 On the PE router, verify the default scan interval.

The PE router (Cisco IOS XR) output should be similar to the following, taken from Pod 3:

```
RP/0/RSP0/CPU0:PE3# show bgp summary
BGP router identifier 10.3.1.1, local AS number 64500
BGP generic scan interval 60 secs
BGP table state: Active
Table ID: 0xe0000000 RD version: 62
BGP main routing table version 62
BGP scan interval 60 secs

BGP is operating in STANDALONE mode.
```

Process	RcvTblVer	bRIB/RIB	LabelVer	ImportVer	SendTblVer	StandbyVer
Speaker	62	62	62	62	62	62

Neighbor	Spk	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	St/PfxRcd
10.0.1.1	0	64500	27291	27285	62	0	0	00:57:54	4
10.0.2.1	0	64500	25985	25980	62	0	0	00:57:56	2
192.168.103.31	0	64503	30009	27293	62	0	0	00:58:25	2

The PE router (Cisco IOS XE) output should be similar to the following, taken from Pod 4:

```

PE4# show ip bgp summary
BGP router identifier 10.4.1.1, local AS number 64500
BGP table version is 44, main routing table version 44
6 network entries using 1488 bytes of memory
8 path entries using 960 bytes of memory
4/4 BGP path/bestpath attribute entries using 960 bytes of memory
1 BGP rrinfo entries using 40 bytes of memory
2 BGP AS-PATH entries using 48 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 3496 total bytes of memory
BGP activity 30/18 prefixes, 47/33 paths, scan interval 60 secs

```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
10.0.1.1	4	64500	27276	29993	44	0	0	2w4d	4
10.0.2.1	4	64500	25923	28523	44	0	0	2w3d	2
192.168.104.41	4	64504	157	161	44	0	0	02:19:00	2

Step 2 On the PE router, set the scan interval to 30 seconds.

Step 3 On the PE router, verify that the scan interval is set to 30 seconds.

The PE router (Cisco IOS XR) output should be similar to the following, taken from Pod 3:

```

RP/0/RSP0/CPU0:PE3# show bgp summary
BGP router identifier 10.3.1.1, local AS number 64500
BGP generic scan interval 30 secs
BGP table state: Active
Table ID: 0xe0000000 RD version: 62
BGP main routing table version 62
BGP scan interval 60 secs

BGP is operating in STANDALONE mode.

```

Process	RcvTblVer	bRIB/RIB	LabelVer	ImportVer	SendTblVer	StandbyVer
Speaker	62	62	62	62	62	62

Neighbor	Spk	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	St/PfxRcd
10.0.1.1	0	64500	27297	27291	62	0	0	01:03:51	4
10.0.2.1	0	64500	25991	25986	62	0	0	01:03:53	2
192.168.103.31	0	64503	30016	27299	62	0	0	01:04:22	2

The PE router (Cisco IOS XE) output should be similar to the following, taken from Pod 4:

```

PE4# show ip bgp summary
BGP router identifier 10.4.1.1, local AS number 64500
BGP table version is 44, main routing table version 44
6 network entries using 1488 bytes of memory
8 path entries using 960 bytes of memory
4/4 BGP path/bestpath attribute entries using 960 bytes of memory
1 BGP rrinfo entries using 40 bytes of memory
2 BGP AS-PATH entries using 48 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 3496 total bytes of memory
BGP activity 30/18 prefixes, 47/33 paths, scan interval 30 secs

```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
10.0.1.1	4	64500	27282	30000	44	0	0	2w4d	4
10.0.2.1	4	64500	25929	28530	44	0	0	2w3d	2
192.168.104.41	4	64504	164	169	44	0	0	02:25:44	2

Step 4 On the CE router, verify the default advertisement interval for the PE neighbor.

The CE router (Cisco IOS) output should be similar to the following, taken from Pod 3:

```

CE3# show ip bgp neighbors 192.168.103.30 | include minimum time
Default minimum time between advertisement runs is 30 seconds

```

The CE router (Cisco IOS) output should be similar to the following, taken from Pod 4:

```

CE4# show ip bgp neighbors 192.168.104.40 | include minimum time
Default minimum time between advertisement runs is 30 seconds

```

Step 5 On the CE router, set the advertisement interval for the PE neighbor to 15 seconds.

Step 6 On the CE router, verify the advertisement interval for the PE neighbor.

The CE router (Cisco IOS) output should be similar to the following, taken from Pod 3:

```

CE3# show ip bgp neighbors 192.168.103.30
BGP neighbor is 192.168.103.30, remote AS 64500, external link
  BGP version 4, remote router ID 10.3.1.1
  BGP state = Established, up for 02:00:37
  Last read 00:00:32, last write 00:00:03, hold time is 180, keepalive interval is 60
seconds
  Neighbor sessions:
    1 active, is not multisession capable (disabled)
  Neighbor capabilities:
    Route refresh: advertised and received(new)
    Four-octets ASN Capability: advertised and received
    Address family IPv4 Unicast: advertised and received
    Enhanced Refresh Capability: advertised
    Multisession Capability:
    Stateful switchover support enabled: NO for session 1
  Message statistics:
    InQ depth is 0
    OutQ depth is 0

      Sent      Rcvd
Opens:          1          1
Notifications:  0          0
Updates:        8          5
Keepalives:    133        121
Route Refresh:  0          0
Total:         142        127
Default minimum time between advertisement runs is 30 seconds

For address family: IPv4 Unicast
Session: 192.168.103.30
BGP table version 41, neighbor version 41/0
Output queue size : 0
Index 7, Advertise bit 0
7 update-group member
Slow-peer detection is disabled
Slow-peer split-update-group dynamic is disabled

      Sent      Rcvd
Prefix activity:  ----
Prefixes Current:    2          4 (Consumes 320 bytes)
Prefixes Total:      5          5
Implicit Withdraw:    0          0
Explicit Withdraw:    3          1
Used as bestpath:     n/a        4
Used as multipath:     n/a        0

      Outbound    Inbound
Local Policy Denied Prefixes:  -----
  Bestpath from this peer:      5          n/a
  Total:                        5          0
Number of NLRI's in the update sent: max 2, min 0
Last detected as dynamic slow peer: never
Dynamic slow peer recovered: never
Minimum time between advertisement runs is 15 seconds
<... output omitted ...>

```

The CE router (Cisco IOS) output should be similar to the following, taken from Pod 4:

```

CE4# show ip bgp neighbors 192.168.104.40
BGP neighbor is 192.168.104.40, remote AS 64500, external link
  BGP version 4, remote router ID 10.4.1.1
  BGP state = Established, up for 03:22:01
  Last read 00:00:27, last write 00:00:41, hold time is 180, keepalive interval is 60
seconds
  Neighbor sessions:
    1 active, is not multisession capable (disabled)
  Neighbor capabilities:
    Route refresh: advertised and received(new)
    Four-octets ASN Capability: advertised and received
    Address family IPv4 Unicast: advertised and received
    Enhanced Refresh Capability: advertised and received
    Multisession Capability:
    Stateful switchover support enabled: NO for session 1
  Message statistics:
    InQ depth is 0
    OutQ depth is 0

              Sent          Rcvd
Opens:             1             1
Notifications:     0             0
Updates:           4            17
Keepalives:       222           216
Route Refresh:     0             0
Total:            227           234
Default minimum time between advertisement runs is 30 seconds

For address family: IPv4 Unicast
Session: 192.168.104.40
BGP table version 45, neighbor version 45/0
Output queue size : 0
Index 5, Advertise bit 0
5 update-group member
Slow-peer detection is disabled
Slow-peer split-update-group dynamic is disabled

              Sent          Rcvd
Prefix activity:  ----
Prefixes Current:      2             4 (Consumes 320 bytes)
Prefixes Total:        3            12
Implicit Withdraw:      0             0
Explicit Withdraw:     1             8
Used as bestpath:      n/a            4
Used as multipath:      n/a            0

              Outbound      Inbound
Local Policy Denied Prefixes:  -----
  Bestpath from this peer:      12          n/a
  Total:                        12           0
Number of NLRI in the update sent: max 2, min 0
Last detected as dynamic slow peer: never
Dynamic slow peer recovered: never
Minimum time between advertisement runs is 15 seconds
Refresh Epoch: 1
<... output omitted ...>

```

Step 7 On the CE router, disable the Loopback0 interface. Disabling the interface will end advertisement of the 10.x.10.1/32 or 10.y.10.1/32 network to the PE router.

After the Loopback0 interface on the CE router is disabled, the PE router (Cisco IOS XR or Cisco IOS XE) output should be similar to the following, taken from Pod 3 and Pod 4, respectively:

```
RP/0/RSP0/CPU0:PE3# show bgp
Tue Jul 29 14:42:17.695 UTC
BGP router identifier 10.3.1.1, local AS number 64500
BGP generic scan interval 30 secs
BGP table state: Active
Table ID: 0xe0000000 RD version: 46
BGP main routing table version 46
BGP scan interval 60 secs
```

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

Network	Next Hop	Metric	LocPrf	Weight	Path
*>i 10.0.1.1/32	10.0.1.1	0	100	0	i
*> 10.3.1.1/32	0.0.0.0	0		32768	i
*> 10.3.100.1/32	192.168.103.31	0		0	64503 i
*>i 10.4.1.1/32	10.4.1.1	0	100	0	i
*>i 10.4.100.1/32	10.4.1.1	0	100	0	64504 i
* i	10.4.1.1	0	100	0	64504 i

Processed 5 prefixes, 6 paths

```
PE4# show ip bgp
```

BGP table version is 44, local router ID is 10.4.1.1

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

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

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

Network	Next Hop	Metric	LocPrf	Weight	Path
r>i 10.0.1.1/32	10.0.1.1	0	100	0	i
r>i 10.3.1.1/32	10.3.1.1	0	100	0	i
*>i 10.3.100.1/32	10.3.1.1	0	100	0	64503 i
* i	10.3.1.1	0	100	0	64503 i
*> 10.4.1.1/32	0.0.0.0	0		32768	i
*> 10.4.100.1/32	192.168.104.41	0		0	64504 i

Immediately bring the interface back up. On the PE router, examine the BGP table. You should see that 15 seconds is required for the network to reappear in the BGP table because the minimum time between successive updates on the CE router is set to 15 seconds.

After the Loopback0 interface on the CE router is enabled, the PE router (Cisco IOS XR or Cisco IOS XE) output should be similar to the following, taken from Pod 3 and Pod 4, respectively:

```
RP/0/RSP0/CPU0:PE3# show bgp
Tue Jul 29 14:43:34.037 UTC
BGP router identifier 10.3.1.1, local AS number 64500
BGP generic scan interval 30 secs
BGP table state: Active
Table ID: 0xe0000000 RD version: 49
BGP main routing table version 49
BGP scan interval 60 secs

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

   Network        Next Hop        Metric LocPrf Weight Path
*>i10.0.1.1/32     10.0.1.1             0      100        0 i
*> 10.3.1.1/32     0.0.0.0              0           32768 i
*> 10.3.10.1/32    192.168.103.31       0           0 64503 i
*> 10.3.100.1/32   192.168.103.31       0           0 64503 i
*>i10.4.1.1/32     10.4.1.1             0      100        0 i
*>i10.4.10.1/32    10.4.1.1             0      100        0 64504 i
* i               10.4.1.1             0      100        0 64504 i
*>i10.4.100.1/32   10.4.1.1             0      100        0 64504 i
* i               10.4.1.1             0      100        0 64504 i

Processed 7 prefixes, 9 paths
```

```
PE4# show ip bgp
BGP table version is 47, local router ID is 10.4.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

   Network        Next Hop        Metric LocPrf Weight Path
r>i 10.0.1.1/32     10.0.1.1             0      100        0 i
r>i 10.3.1.1/32     10.3.1.1             0      100        0 i
*>i 10.3.10.1/32    10.3.1.1             0      100        0 64503 i
* i               10.3.1.1             0      100        0 64503 i
*>i 10.3.100.1/32   10.3.1.1             0      100        0 64503 i
* i               10.3.1.1             0      100        0 64503 i
*> 10.4.1.1/32     0.0.0.0              0           32768 i
*> 10.4.10.1/32    192.168.104.41       0           0 64504 i
*> 10.4.100.1/32   192.168.104.41       0           0 64504 i
```

Verification

No verification is needed for this task.

Task 4: Implement BGP Route Dampening

In this task, you will implement BGP route dampening.

Activity Procedure

Complete these steps:

- Step 1** On the PE router, enable BGP route dampening with the default parameters. Enable debugging of BGP dampening events.
- Step 2** On the CE router, disable the Loopback0 interface. Disabling the interface will end advertisement of the 10.x.10.1/32 or 10.y.10.1/32 network to the PE router.
- Step 3** On the PE router, verify the BGP table. You should see that the route 10.x.10.1/32 or 10.y.10.1/32 was not removed from the table.

The PE router (Cisco IOS XR) output should be similar to the following, taken from Pod 3:

```
RP/0/RSP0/CPU0:PE3# show bgp
Tue Jul 29 14:46:47.232 UTC
BGP router identifier 10.3.1.1, local AS number 64500
BGP generic scan interval 30 secs
BGP table state: Active
Table ID: 0xe0000000 RD version: 51
BGP main routing table version 51
Dampening enabled
BGP scan interval 60 secs

Status codes: s suppressed, d damped, h history, * valid, > best
                i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network                Next Hop              Metric LocPrf Weight Path
*>i10.0.1.1/32            10.0.1.1                0      100      0 i
*>10.3.1.1/32             0.0.0.0                  0          32768 i
h 10.3.10.1/32            192.168.103.31           0          0 64503 i
*>10.3.100.1/32          192.168.103.31           0          0 64503 i
*>i10.4.1.1/32            10.4.1.1                0      100      0 i
*>i10.4.100.1/32         10.4.1.1                0      100      0 64504 i
* i                       10.4.1.1                0      100      0 64504 i

Processed 6 prefixes, 7 paths
```

The PE router (Cisco IOS XE) output should be similar to the following, taken from Pod 4:

```
PE4# show ip bgp
BGP table version is 49, local router ID is 10.4.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
                r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
                x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

   Network                Next Hop              Metric LocPrf Weight Path
r>i 10.0.1.1/32            10.0.1.1                0      100      0 i
r>i 10.3.1.1/32            10.3.1.1                0      100      0 i
*>i 10.3.100.1/32          10.3.1.1                0      100      0 64503 i
* i                       10.3.1.1                0      100      0 64503 i
*> 10.4.1.1/32             0.0.0.0                  0          32768 i
h 10.4.10.1/32            192.168.104.41           0          0 64504 i
*> 10.4.100.1/32          192.168.104.41           0          0 64504 i
```

- Step 4** On the CE router, enable the Loopback0 interface to come back up. Disable and enable the interface a few times. Wait for more than 15 seconds between disabling and enabling the interface because the advertisement interval is set to 15 seconds. Disabling and enabling the interface a few times causes the 10.x.10.1/32 or 10.y.10.1/32 network to flap.
- Step 5** On the PE router, watch the logging messages about penalizing and eventually suppressing the route.

The PE router (Cisco IOS XR) output should be similar to the following, taken from Pod 3:

```
RP/0/RSP0/CPU0:Jun 11 13:40:15.074 : bgp[1049]: [default-rtr] (ip4u): Charge penalty
for 10.3.10.1/32 path 64503 with halflife-time 15 min reuse/suppress 750/2000 Flapped 4
times in 00:05:55. New penalty is 3510
RP/0/RSP0/CPU0:Jun 11 13:40:15.085 : bgp[1049]: [default-iowt]: Keepalive timer started
for 2001:db8:10:0:1::1(loc 10): last 3800819 this 3800832
```

The PE router (Cisco IOS XE) output should be similar to the following, taken from Pod 4:

```
*Jun 11 21:22:18.114: BGP(0): suppress 10.4.10.1/32 path 64504 for 00:28:40 (penalty
2819)
*Jun 11 21:22:18.114: halflife-time 15, reuse/suppress 750/2000
*Jun 11 21:22:18.114: EvD: accum. penalty 2819, now suppressed with a reuse intervals
of 172
```

- Step 6** How many times did you have to flap the route for the PE router to suppress the route?

- Step 7** What is the default suppress penalty, the half-life time, and the reuse penalty?

- Step 8** On the PE router, disable BGP dampening debugging.

Verification

You have completed this task when you attain these results:

- Step 1** On the PE router, verify the BGP table. You should see that the route was not removed from the table. The route should be put into the history or dampened state.

The PE router (Cisco IOS XR) output should be similar to the following, taken from Pod 3:

```

RP/0/RSP0/CPU0:PE3# show bgp
Wed Jun 11 13:42:32.338 UTC
BGP router identifier 10.3.1.1, local AS number 64500
BGP generic scan interval 30 secs
BGP table state: Active
Table ID: 0xe0000000 RD version: 1334
BGP main routing table version 1334
Dampening enabled
BGP scan interval 60 secs

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

   Network        Next Hop        Metric LocPrf Weight Path
*>i10.0.1.1/32     10.0.1.1             0      100      0 i
*> 10.3.1.1/32     0.0.0.0              0           32768 i
*d 10.3.10.1/32    192.168.103.31       0           0 64503 i
*> 10.3.100.1/32   192.168.103.31       0           0 64503 i
*>i10.4.1.1/32     10.4.1.1             0      100      0 i
* i               10.4.1.1             0      100      0 i
*>i10.4.100.1/32   10.4.1.1             0      100      0 64504 i
* i               10.4.1.1             0      100      0 64504 i

Processed 5 prefixes, 7 paths

```

The PE router (Cisco IOS XE) output should be similar to the following, taken from Pod 4:

```

PE4# show ip bgp
BGP table version is 19, local router ID is 10.4.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
               x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

   Network        Next Hop        Metric LocPrf Weight Path
r>i 10.0.1.1/32     10.0.1.1             0      100      0 i
r>i 10.3.1.1/32     10.3.1.1             0      100      0 i
* i 10.3.100.1/32   10.3.1.1             0      100      0 64503 i
*>i               10.3.1.1             0      100      0 64503 i
*> 10.4.1.1/32     0.0.0.0              0           32768 i
*d 10.4.10.1/32    192.168.104.41       0           0 64504 i
*> 10.4.100.1/32   192.168.104.41       0           0 64504 i

```

Step 2 On the PE router, verify the dampened routes.

The PE router (Cisco IOS XR) output should be similar to the following, taken from Pod 3:

```
RP/0/RSP0/CPU0:PE3# show bgp ipv4 unicast dampened-paths
Wed Jun 11 13:44:58.037 UTC
BGP router identifier 10.3.1.1, local AS number 64500
BGP generic scan interval 30 secs
BGP table state: Active
Table ID: 0xe0000000 RD version: 1334
BGP main routing table version 1334
Dampening enabled
BGP scan interval 60 secs

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

   Network          From          Reuse      Path
*d 10.3.10.1/32     192.168.103.31 00:28:40 64503 i
```

The PE router (Cisco IOS XE) output should be similar to the following, taken from Pod 4:

```
PE4# show ip bgp dampening dampened-paths
BGP table version is 19, local router ID is 10.4.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
               x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

   Network          From          Reuse      Path
*d 10.4.10.1/32     192.168.104.41 00:09:54 64504 I
```

Step 3 On the PE router, examine information about the 10.x.10.1/32 or 10.y.10.1/32 route. You should see BGP dampening information about the route.

The PE router (Cisco IOS XR) output should be similar to the following, taken from Pod 3:

```
RP/0/RSP0/CPU0:PE3# show bgp 10.3.10.1/32
Wed Jun 11 13:46:31.917 UTC
BGP routing table entry for 10.3.10.1/32
Versions:
  Process          bRIB/RIB  SendTblVer
  Speaker          1333     1333
Last Modified: Jun 11 13:37:31.829 for 00:09:00
Paths: (1 available, no best path)
  Not advertised to any peer
  Path #1: Received by speaker 0
  Not advertised to any peer
  64503, (suppressed due to dampening)
    192.168.103.31 from 192.168.103.31 (10.3.100.1)
      Origin IGP, metric 0, localpref 100, valid, external, import suspect
      Received Path ID 0, Local Path ID 0, version 0
      Origin-AS validity: not-found
      Dampinfo: penalty 2625, flapped 4 times in 00:12:12, reuse in 00:27:00
                half life 00:15:00, suppress value 2000, reuse value 750
                Maximum suppress time 01:00:00
```

The PE router (Cisco IOS XE) output should be similar to the following, taken from Pod 4:

```
PE4# show ip bgp 10.4.10.1/32
BGP routing table entry for 10.4.10.1/32, version 19
Paths: (1 available, no best path)
  Not advertised to any peer
  Refresh Epoch 1
  64504, (suppressed due to dampening)
    192.168.104.41 from 192.168.104.41 (10.4.100.1)
      Origin IGP, metric 0, localpref 100, valid, external
      Dampinfo: penalty 1842, flapped 3 times in 00:12:30, reuse in 00:08:34
      rx pathid: 0, tx pathid: 0
```

Hardware Lab 4: Implement Layer 2 and Layer 3 Multicast

Activity Objective

In this lab activity, you will implement and verify the operation of IGMP and MLD and you will observe multicast flooding on the LAN when IGMP snooping is implemented.


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

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

- Configure IP multicast support and monitor for IGMP and MLD
- Monitor the network without and with IGMP snooping

Visual Objective

The figure illustrates what you will accomplish in this activity.



Command List

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

Cisco IOS/IOS XE Software Commands

Command	Description
[no] ip igmp snooping	Enables or disables IGMP snooping globally on the switch.
configure terminal	Enters configuration mode.
interface <i>interface</i>	Enters interface configuration mode.
ip igmp join-group <i>group</i>	Enables an interface to join a multicast group.
ip igmp query-interval <i>interval</i>	Sets the IGMP query interval on the interface.
ip igmp version <i>version</i>	Sets the IGMP router version on the interface.
ip multicast-routing distributed	Enables IP multicast routing globally on the router.
ipv6 mld query-interval <i>interval</i>	Sets the MLD query interval on the interface.

Command	Description
ipv6 mld router	Enables the MLD router on the interface.
ping <i>dest_ip source source_interface</i>	Verifies connectivity between the source IP and destination IP.
show [ip ipv6] [igmp mld] groups <i>interface</i>	Displays IGMP or MLD group information.
show [ip ipv6] [igmp mld] interface <i>interface</i>	Displays IGMP or MLD interface information.
show ip igmp snooping groups	Displays IGMP snooping information on the switch.

Cisco IOS XR Software Commands

Command	Description
address-family [ipv4 ipv6]	Enters an IPv4 or IPv6 address family in multicast routing mode.
commit	Commits changes to the running configuration.
configure terminal	Enters configuration mode.
enable	Enables multicast routing on the interface.
interface <i>interface</i>	Enters interface configuration mode.
multicast-routing	Enables multicast routing and enters multicast routing configuration mode.
ping <i>dest_IP source source_IP</i>	Verifies connectivity between the source IP and destination IP (IPv4 and IPv6).
query-interval <i>interval</i>	Sets the IGMP or MLD query interval.
router enable	Enables the MLD router.
router mld	Enables the MLD router and enters MLD router configuration mode.
show [igmp mld] groups <i>interface</i>	Displays IGMP or MLD group information.
show [igmp mld] interface <i>interface</i>	Displays IGMP or MLD interface information.
version <i>version</i>	Sets the IGMP or MLD router version.

Task 1: Enable IGMP and MLD

In this task, you will configure and verify IGMP and MLD support on the pod PE router.

Activity Procedure

Complete these steps:

Step 1 On the pod PE router, enable IPv4 and IPv6 multicast routing on the first Gigabit Ethernet interface.

Step 2 On the pod PE router, verify the IGMP version and query interval:

The PE router (Cisco IOS XR) output should be similar to the following, taken from Pod 3:

```
RP/0/RSP0/CPU0:PE3# show igmp interface GigabitEthernet 0/0/0/0
GigabitEthernet0/0/0/0 is up, line protocol is up
  Internet address is 192.168.103.30/24
  IGMP is enabled on interface
  Current IGMP version is 3
  IGMP query interval is 60 seconds
  IGMP querier timeout is 125 seconds
  IGMP max query response time is 10 seconds
  Last member query response interval is 1 seconds
  IGMP activity: 4 joins, 0 leaves
  IGMP querying router is 192.168.103.30 (this system)
  Time elapsed since last query sent 00:00:49
  Time elapsed since IGMP router enabled 00:01:01
  Time elapsed since last report received 00:00:45
```

The PE router (Cisco IOS XE) output should be similar to the following, taken from Pod 4:

```
PE4# show ip igmp interface GigabitEthernet 0/0/0
GigabitEthernet0/0/0 is up, line protocol is up
  Internet address is 192.168.104.40/24
  IGMP is disabled on interface
  Multicast routing is disabled on interface
  Multicast TTL threshold is 0
  No multicast groups joined by this system
```

IGMP is disabled on the PE router running Cisco IOS XE Software. IGMP is automatically enabled when PIM is configured on the interface. Enable PIM sparse mode on the Gigabit Ethernet interface by using the following commands:

```
PEy(config)#
interface GigabitEthernet 0/0/0
 ip pim sparse-mode
```

y should be substituted with 2 or 4, depending on your team assignment.

Step 3 On the pod PE router, verify the MLD version and query interval.

The PE router (Cisco IOS XR) output should be similar to the following, taken from Pod 3:

```
RP/0/RSP0/CP00:PE3# show mld interface GigabitEthernet 0/0/0/0
GigabitEthernet0/0/0/0 is up, line protocol is up
Internet address is fe80::6e9c:edff:fe22:55c8
MLD is enabled on interface
Current MLD version is 2
MLD query interval is 125 seconds
MLD querier timeout is 255 seconds
MLD max query response time is 10 seconds
Last member query response interval is 1 seconds
MLD activity: 5 joins, 0 leaves
MLD querying router is fe80::6e9c:edff:fe22:55c8 (this system)
Time elapsed since last query sent 00:00:03
Time elapsed since IGMP router enabled 00:06:40
Time elapsed since last report received 00:00:01
```

The PE router (Cisco IOS XE) output should be similar to the following, taken from Pod 4:

```
PE4# show ipv6 mld interface GigabitEthernet 0/0/0
GigabitEthernet0/0/0 is up, line protocol is up
Internet address is FE80::462B:3FF:FE53:FC00/10
MLD is enabled on interface
Current MLD version is 2
MLD query interval is 125 seconds
MLD querier timeout is 255 seconds
MLD max query response time is 10 seconds
Last member query response interval is 1 seconds
MLD activity: 7 joins, 1 leaves
MLD querying router is FE80::462B:3FF:FE53:FC00 (this system)
```

MLD is automatically enabled when IPv6 multicast routing is enabled.

- Step 4** On the pod PE router of the first GigabitEthernet interface, configure IGMP version 2 and the IGMP query interval at 30 seconds.
- Step 5** On the pod PE router of the first Gigabit Ethernet interface, enable the MLD router and configure the MLD query interval at 60 seconds.
- Step 6** On the pod CE router, configure the first Gigabit Ethernet interface to join multicast group 234.1.1.1.

Verification

You have completed this task when you attain these results:

- Step 1** On the pod PE router, verify the IGMP version and query interval. The IGMP version should be 2 and query interval should be 30 seconds.

The PE router (Cisco IOS XR) output should be similar to the following, taken from Pod 3:

```
RP/0/RSP0/CPU0:PE3# show igmp interface GigabitEthernet 0/0/0/0
GigabitEthernet0/0/0/0 is up, line protocol is up
Internet address is 192.168.103.30/24
IGMP is enabled on interface
Current IGMP version is 2
IGMP query interval is 30 seconds
IGMP querier timeout is 65 seconds
IGMP max query response time is 10 seconds
Last member query response interval is 1 seconds
IGMP activity: 5 joins, 1 leaves
IGMP querying router is 192.168.103.30 (this system)
Time elapsed since last query sent 00:00:25
Time elapsed since IGMP router enabled 23:06:51
Time elapsed since last report received 00:00:22
```

The PE router (Cisco IOS XE) output should be similar to the following, taken from Pod 4:

```
PE4# show ip igmp interface GigabitEthernet 0/0/0
GigabitEthernet0/0/0 is up, line protocol is up
Internet address is 192.168.104.40/24
IGMP is enabled on interface
Current IGMP host version is 2
Current IGMP router version is 2
IGMP query interval is 30 seconds
IGMP configured query interval is 30 seconds
IGMP querier timeout is 60 seconds
IGMP configured querier timeout is 60 seconds
IGMP max query response time is 10 seconds
Last member query count is 2
Last member query response interval is 1000 ms
Inbound IGMP access group is not set
IGMP activity: 2 joins, 0 leaves
Multicast routing is enabled on interface
Multicast TTL threshold is 0
Multicast designated router (DR) is 192.168.104.40 (this system)
IGMP querying router is 192.168.104.40 (this system)
Multicast groups joined by this system (number of users):
  224.0.1.40(1)
```

Step 2 On the pod PE router, verify the MLD version and query interval. The MLD version should be 2 and query interval should be 60 seconds.

The PE router (Cisco IOS XR) output should be similar to the following, taken from Pod 3:

```
RP/0/RSP0/CPU0:PE3# show mld interface GigabitEthernet 0/0/0/0
GigabitEthernet0/0/0/0 is up, line protocol is up
Internet address is fe80::6e9c:edff:fe22:55c8
MLD is enabled on interface
Current MLD version is 2
MLD query interval is 60 seconds
MLD querier timeout is 125 seconds
MLD max query response time is 10 seconds
Last member query response interval is 1 seconds
MLD activity: 5 joins, 0 leaves
MLD querying router is fe80::6e9c:edff:fe22:55c8 (this system)
Time elapsed since last query sent 00:00:05
Time elapsed since IGMP router enabled 23:52:28
Time elapsed since last report received 00:00:04
```

The PE router (Cisco IOS XE) output should be similar to the following, taken from Pod 4:

```
PE4# show ipv6 mld interface GigabitEthernet 0/0/0
GigabitEthernet0/0/0 is up, line protocol is up
Internet address is FE80::462B:3FF:FE53:FC00/10
MLD is enabled on interface
Current MLD version is 2
MLD query interval is 60 seconds
MLD querier timeout is 125 seconds
MLD max query response time is 10 seconds
Last member query response interval is 1 seconds
MLD activity: 7 joins, 1 leaves
MLD querying router is FE80::462B:3FF:FE53:FC00 (this system)
```

Step 3 On the pod PE router, verify IGMP groups on the first Gigabit Ethernet interface. Verify that multicast group 234.1.1.1 appears in the IGMP groups table.

The PE router (Cisco IOS XR) output should be similar to the following, taken from Pod 3:

```
RP/0/RSP0/CPU0:PE3# show igmp groups GigabitEthernet 0/0/0/0
Fri Jun 13 07:30:25.510 UTC
IGMP Connected Group Membership
Group Address Interface Uptime Expires Last Reporter
224.0.0.2 GigabitEthernet0/0/0/0 23:34:26 never 192.168.103.30
224.0.0.13 GigabitEthernet0/0/0/0 23:34:26 never 192.168.103.30
224.0.1.40 GigabitEthernet0/0/0/0 23:34:26 never 192.168.103.30
234.1.1.1 GigabitEthernet0/0/0/0 00:30:46 00:00:46 192.168.103.31
```

The PE router (Cisco IOS XE) output should be similar to the following, taken from Pod 4:

```
PE4# show ip igmp groups GigabitEthernet 0/0/0
IGMP Connected Group Membership
Group Address Interface Uptime Expires Last Reporter Group
Accounted
234.1.1.1 GigabitEthernet0/0/0 00:19:07 00:01:24 192.168.104.41
224.0.1.40 GigabitEthernet0/0/0 00:19:10 00:01:22 192.168.104.40
```

Task 2: Verify IGMP Snooping

In this task, you will verify IGMP snooping on the pod switch. You will disable IGMP snooping and verify the results.

Activity Procedure

Complete these steps:

Step 1 On the pod switch, verify IGMP snooping:

The SW switch output should be similar to the following, taken from Pod 3:

```
SW3# show ip igmp snooping groups
```

Vlan	Group	Type	Version	Port List
1	224.0.0.1:40	igmp	v2,v3	Fa0/2
1	234.1.1.1	igmp	v2	Fa0/1, Fa0/2

The SW switch output should be similar to the following, taken from Pod 4:

```
SW4# show ip igmp snooping groups
```

Vlan	Group	Type	Version	Port List
1	224.0.0.1:40	igmp	v2	Fa0/2
1	234.1.1.1	igmp	v2	Fa0/1, Fa0/2

Step 2 On the pod switch, disable IGMP snooping.

Step 3 On the pod switch, verify that IGMP snooping is disabled.

The SW switch output should be similar to the following, taken from Pod 3:

```
SW3# show ip igmp snooping groups
SW3#
```

The SW switch output should be similar to the following, taken from Pod 4:

```
SW4# show ip igmp snooping groups
SW4#
```

Step 4 On the pod switch, enable IGMP snooping.

Step 5 On the pod CE router, configure the first Gigabit Ethernet interface to leave multicast group 234.1.1.1.

Verification

No verification is needed for this task.

Hardware Lab 5: Enable and Optimize PIM-SM

Activity Objective

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

In this lab activity, you will enable multicast on the router in your team. The P1 router is preconfigured to act as an RP for your multicast traffic. You will configure receivers for multicast traffic on the CE and PE router. The other pod CE router will act as a multicast source.


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

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

- Implement multicast routing, PIM-SM, and manual RP configuration
- Observe shared tree formation
- Observe the switchover from the shared tree to the SPT

Visual Objective

The figure illustrates what you will accomplish in this activity.



Command List

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

Cisco IOS/IOS XE Software Commands

Command	Description
configure terminal	Enters configuration mode.
interface interface	Enters interface configuration mode.
ip igmp join-group group_address	Configures an interface on the router to join the specified group or channel.
ip multicast-routing	Enables IP multicast routing.
ip pim rp-address RP_address	Statically configures the address of a PIM RP for multicast groups.
ip pim sparse-mode	Enables an interface for PIM-SM.
ip pim spt-threshold threshold	Configures when a PIM leaf router should join the shortest path source tree.

Command	Description
ping <i>dest_IP source source_IP</i>	Verifies connectivity between source IP and destination IP (IPv4 and IPv6).
show ip mroute	Displays the contents of the multicast routing table.
show ip pim interface	Displays information about interfaces that are configured for PIM.
show ip pim neighbor	Displays information about PIM neighbors discovered by PIMv1 router query messages or PIMv2 hello messages.
show ip route	Displays the routing table on a router.

Cisco IOS XR Software Commands

Command	Description
show route	Displays the routing table on a router.
address-family ipv4	Enters an IPv4 address family under a specific configuration mode.
commit	Commits changes to the running configuration.
configure	Enters configuration mode.
show mrib route	Displays the contents of the multicast routing table.
show pim neighbor	Displays information about PIM neighbors discovered by PIMv1 router query messages or PIMv2 hello messages.
show pim interface	Displays information about interfaces that are configured for PIM.
interface <i>interface</i>	Enters interface configuration mode.
enable	Enables an interface for multicast routing or PIM (under the appropriate configuration mode).
multicast-routing	Enters multicast routing configuration mode.
router pim	Enters PIM configuration mode.
ping <i>dest_IP source source_IP</i>	Verifies connectivity between the source IP and the destination IP (IPv4 and IPv6).
rp-address <i>RP_address</i>	Statically configures the address of a PIM RP for multicast groups under PIM configuration mode.

Command	Description
spt-threshold infinity	Configures a PIM leaf router to join the shortest path source tree immediately.
ip igmp join-group <i>group_address</i>	Configures an interface on the router to join the specified group or channel..
router igmp	Enters IGMP configuration mode.

Task 1: Implement PIM-SM

In this task, you will enable multicast on all routers.

Activity Procedure

Complete these steps:

- Step 1** On the PE router make sure that the route for the P1 Loopback0 (10.0.1.1/32) interface points to GigabitEthernet0/0/0/2.

The PE router (Cisco IOS XR) output should be similar to the following, taken from Pod 3:

```
RP/0/RSP0/CPU0:PE3# show route 10.0.1.1
Routing entry for 10.0.1.1/32
  Known via "isis 1", distance 115, metric 20, type level-2
  Installed May 22 13:27:23.304 for 3w6d
  Routing Descriptor Blocks
    192.168.31.1, from 10.0.1.1, via GigabitEthernet0/0/0/2
    Route metric is 20
  No advertising protos.
```

The PE router (Cisco IOS XE) output should be similar to the following, taken from Pod 4:

```
PE4# show ip route 10.0.1.1
Routing entry for 10.0.1.1/32
  Known via "isis", distance 115, metric 20, type level-2
  Redistributing via isis
  Last update from 192.168.41.1 on GigabitEthernet0/0/2, 1w0d ago
  Routing Descriptor Blocks:
    * 192.168.41.1, from 10.0.1.1, 1w0d ago, via GigabitEthernet0/0/2
    Route metric is 20, traffic share count is 1
```

- Step 2** On the PE router make sure that the route for the other pod CE router Loopback0 (10.x.10.1/32 or 10.y.10.1/32) interface points to GigabitEthernet0/0/0/1.

The PE router (Cisco IOS XR) output should be similar to the following, taken from Pod 3:

```

RP/0/RSP0/CPU0:PE3# show route 10.4.10.1
Routing entry for 10.4.10.1/32
  Known via "bgp 64500", distance 200, metric 0
  Tag 64504, type internal
  Installed Jun 11 14:06:44.009 for 1w0d
  Routing Descriptor Blocks
    10.4.1.1, from 10.0.1.1
      Route metric is 0
  No advertising protos.

RP/0/RSP0/CPU0:PE3# show route 10.4.1.1
Routing entry for 10.4.1.1/32
  Known via "isis 1", distance 115, metric 20, type level-2
  Installed Jun 11 09:11:03.322 for 1w0d
  Routing Descriptor Blocks
    192.168.134.40, from 10.4.1.1, via GigabitEthernet0/0/0/1
      Route metric is 20
  No advertising protos.

```

The PE router (Cisco IOS XE) output should be similar to the following, taken from Pod 4:

```

PE4# show ip route 10.3.10.1
Routing entry for 10.3.10.1/32
  Known via "bgp 64500", distance 200, metric 0
  Tag 64503, type internal
  Last update from 10.3.1.1 1w0d ago
  Routing Descriptor Blocks:
  * 10.3.1.1, from 10.0.1.1, 1w0d ago
    Route metric is 0, traffic share count is 1
    AS Hops 1
    Route tag 64503
    MPLS label: none

PE4# show ip route 10.3.1.1
Routing entry for 10.3.1.1/32
  Known via "isis", distance 115, metric 20, type level-2
  Redistributing via isis
  Last update from 192.168.134.30 on GigabitEthernet0/0/1, 1w0d ago
  Routing Descriptor Blocks:
  * 192.168.134.30, from 10.3.1.1, 1w0d ago, via GigabitEthernet0/0/1
    Route metric is 20, traffic share count is 1

```

- Step 3** Enable IP multicast routing on the PE and CE routers. Enable multicast on all interfaces that have IP addresses assigned (Cisco IOS XR).
- Step 4** Enable PIM-SM on the CE and PE routers on all interfaces that have IP addresses assigned.
- Step 5** On the CE and PE router, define the SPT threshold as infinity. This action should force the routers to always stay on the shared tree.
- Step 6** Manually configure the RP address on the CE and PE router. The P1 router with Loopback0 (10.0.1.1) will act as the RP for all multicast groups.

Verification

Both pods from the same team should be finished with the previous steps at this point.

You have completed this task when you attain these results:

Step 1 Verify PIM state on interfaces on the PE router.

The PE router (Cisco IOS XR) output should be similar to the following, taken from Pod 3:

```
RP/0/RSP0/CPU0:PE3# show pim interface
```

PIM interfaces in VRF default

Address	Interface	PIM	Nbr Count	Hello Intvl	DR Prior	DR
10.3.1.1	Loopback0	on	1	30	1	this system
192.168.103.30	GigabitEthernet0/0/0/0	on	2	30	1	
192.168.103.31						
192.168.134.30	GigabitEthernet0/0/0/1	on	2	30	1	
192.168.134.40						
192.168.31.30	GigabitEthernet0/0/0/2	on	2	30	1	this system
192.168.32.30	GigabitEthernet0/0/0/3	on	2	30	1	this system

The PE router (Cisco IOS XE) output should be similar to the following, taken from Pod 4:

```
PE4# show ip pim interface
```

Address	Interface	Ver/ Mode	Nbr Count	Query Intvl	DR Prior	DR
192.168.104.40	GigabitEthernet0/0/0	v2/S	1	30	1	192.168.104.41
10.4.1.1	Loopback0	v2/S	0	30	1	10.4.1.1
192.168.134.40	GigabitEthernet0/0/1	v2/S	1	30	1	192.168.134.40
192.168.41.40	GigabitEthernet0/0/2	v2/S	1	30	1	192.168.41.40
192.168.42.40	GigabitEthernet0/0/3	v2/S	1	30	1	192.168.42.40

Step 2 Verify PIM neighbors on the PE router. The PE router output should be similar to the following, taken from Pod 1:

The PE router (Cisco IOS XR) output should be similar to the following, taken from Pod 3:

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

PIM neighbors in VRF default

Flag: B - Bidir capable, P - Proxy capable, DR - Designated Router,

E - ECMP Redirect capable

* indicates the neighbor created for this router

Neighbor Address	Interface	Uptime	Expires	DR pri	Flags
10.3.1.1*	Loopback0	02:13:45	00:01:36	1 (DR)	B P E
192.168.103.30*	GigabitEthernet0/0/0/0	1w0d	00:01:40	1	B P E
192.168.103.31	GigabitEthernet0/0/0/0	01:48:32	00:01:28	1 (DR)	P
192.168.134.30*	GigabitEthernet0/0/0/1	02:13:45	00:01:41	1	B P E
192.168.134.40	GigabitEthernet0/0/0/1	01:49:00	00:01:39	1 (DR)	P
192.168.31.1	GigabitEthernet0/0/0/2	02:13:39	00:01:42	1	B P
192.168.31.30*	GigabitEthernet0/0/0/2	02:13:45	00:01:26	1 (DR)	B P E
192.168.32.2	GigabitEthernet0/0/0/3	02:13:36	00:01:29	1	B P
192.168.32.30*	GigabitEthernet0/0/0/3	02:13:45	00:01:16	1 (DR)	B P E

The PE router (Cisco IOS XE) output should be similar to the following, taken from Pod 4:

```

PE4# show ip pim neighbor
PIM Neighbor Table
Mode: B - Bidir Capable, DR - Designated Router, N - Default DR Priority,
      P - Proxy Capable, S - State Refresh Capable, G - GenID Capable
Neighbor      Interface      Uptime/Expires    Ver  DR
Address
192.168.104.41 GigabitEthernet0/0/0  01:49:43/00:01:18 v2   1 / DR S P G
192.168.134.30 GigabitEthernet0/0/1  01:50:11/00:01:26 v2   1 / P G
192.168.41.1   GigabitEthernet0/0/2  01:50:13/00:01:38 v2   1 / P G
192.168.42.2   GigabitEthernet0/0/3  01:50:12/00:01:25 v2   1 / P G

```

Task 2: Shared Tree Formation—Receivers

In this task, you will configure multicast receivers. You will observe multicast routing tables on routers when receivers announce their presence.

Activity Procedure

Complete these steps:

- Step 1** Simulate multicast receivers for group 224.1.x.1 (or 224.1.y.1) on the Loopback0 interface on the CE and PE routers.

NOTE: Throughout the lab exercise, use the **ip igmp join-group** command on Cisco IOS and IOS-XE Software under interface configuration mode to simulate multicast receivers. On the Cisco IOS XR Software, use the **join-group** command under interface configuration mode under router IGMP configuration mode.

- Step 2** Examine the multicast routing table on the PE, CE, and PI routers.

NOTE: Use Telnet to connect to the PI router to examine the multicast routing table.

The CE router (Cisco IOS) output should be similar to the following, taken from Pod 3:

```

CE3# show ip mroute
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector
Outgoing interface flags: H - Hardware switched, A - Assert winner
Timers: Uptime/Expires
Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 224.1.3.1), 20:26:15/00:02:57, RP 10.0.1.1, flags: SCL
  Incoming interface: GigabitEthernet0/0, RPF nbr 192.168.103.30
  Outgoing interface list:
    Loopback0, Forward/Sparse, 20:26:15/00:02:57

(*, 224.0.1.40), 22:23:23/00:02:54, RP 10.0.1.1, flags: SCL
  Incoming interface: GigabitEthernet0/0, RPF nbr 192.168.103.30
  Outgoing interface list:
    Loopback0, Forward/Sparse, 22:23:21/00:02:54

```

The PE router (Cisco IOS XR) output should be similar to the following, taken from Pod 3:

```
RP/0/RSP0/CPU0:PE3# show mrib route
IP Multicast Routing Information Base
Entry flags: L - Domain-Local Source, E - External Source to the Domain,
C - Directly-Connected Check, S - Signal, IA - Inherit Accept,
IF - Inherit From, D - Drop, ME - MDT Encap, EID - Encap ID,
MD - MDT Decap, MT - MDT Threshold Crossed, MH - MDT interface handle
CD - Conditional Decap, MPLS - MPLS Decap, MF - MPLS Encap, EX - Extranet
MoFE - MoFRR Enabled, MoFS - MoFRR State, MoFP - MoFRR Primary
MoFB - MoFRR Backup, RPFID - RPF ID Set
Interface flags: F - Forward, A - Accept, IC - Internal Copy,
NS - Negate Signal, DP - Don't Preserve, SP - Signal Present,
II - Internal Interest, ID - Internal Disinterest, LI - Local Interest,
LD - Local Disinterest, DI - Decapsulation Interface
EI - Encapsulation Interface, MI - MDT Interface, LVIF - MPLS Encap,
EX - Extranet, A2 - Secondary Accept, MT - MDT Threshold Crossed,
MA - Data MDT Assigned, LMI - mLDP MDT Interface, TMI - P2MP-TE MDT Interface
IRMI - IR MDT Interface

(*,224.0.0.0/4) RPF nbr: 192.168.31.1 Flags: C P
  Up: 21:40:45

(*,224.0.0.0/24) Flags: D P
  Up: 1w0d

(*,224.0.1.39) Flags: S P
  Up: 1w0d

(*,224.0.1.40) Flags: S P
  Up: 1w0d
  Outgoing Interface List
    GigabitEthernet0/0/0/0 Flags: II LI, Up: 1w0d

(*,224.1.3.1) RPF nbr: 192.168.31.1 Flags: C
  Up: 20:28:23
  Incoming Interface List
    GigabitEthernet0/0/0/2 Flags: A, Up: 20:28:23
  Outgoing Interface List
    Loopback0 Flags: F IC NS II LI, Up: 20:28:23
    GigabitEthernet0/0/0/0 Flags: F NS, Up: 18:30:16

(*,232.0.0.0/8) Flags: D P
  Up: 1w0d
```

The CE router (Cisco IOS) output should be similar to the following, taken from Pod 4:

```

CE4# show ip mroute
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector
Outgoing interface flags: H - Hardware switched, A - Assert winner
Timers: Uptime/Expires
Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 224.1.1.4.1), 20:28:05/00:02:37, RP 10.0.1.1, flags: SCL
  Incoming interface: GigabitEthernet0/0, RPF nbr 192.168.104.40
  Outgoing interface list:
    Loopback0, Forward/Sparse, 20:28:05/00:02:24

(*, 224.0.1.40), 22:25:16/00:02:59, RP 10.0.1.1, flags: SCL
  Incoming interface: GigabitEthernet0/0, RPF nbr 192.168.104.40
  Outgoing interface list:
    Loopback0, Forward/Sparse, 22:25:14/00:02:32

```

The PE router (Cisco IOS XE) output should be similar to the following, taken from Pod 4:

```

PE4# show ip mroute
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
Timers: Uptime/Expires
Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 224.1.1.4.1), 20:30:09/00:02:51, RP 10.0.1.1, flags: SCL
  Incoming interface: GigabitEthernet0/0/2, RPF nbr 192.168.41.1
  Outgoing interface list:
    GigabitEthernet0/0/0, Forward/Sparse, 18:32:39/00:02:51
    Loopback0, Forward/Sparse, 20:30:09/00:01:59

(192.168.104.41, 224.1.1.4.1), 00:01:37/00:01:22, flags: LT
  Incoming interface: GigabitEthernet0/0/0, RPF nbr 192.168.104.41
  Outgoing interface list:
    Loopback0, Forward/Sparse, 00:01:37/00:01:59

(*, 224.0.1.40), 1w0d/00:02:52, RP 10.0.1.1, flags: SCL
  Incoming interface: GigabitEthernet0/0/2, RPF nbr 192.168.41.1
  Outgoing interface list:
    GigabitEthernet0/0/0, Forward/Sparse, 18:32:39/00:02:37

```

The P1 router (Cisco IOS XR) output should be similar to the following:

```
RP/0/RSP0/CPU0:P1# show mrib route
IP Multicast Routing Information Base
Entry flags: L - Domain-Local Source, E - External Source to the Domain,
             C - Directly-Connected Check, S - Signal, IA - Inherit Accept,
             IF - Inherit From, D - Drop, ME - MDT Encap, EID - Encap ID,
             MD - MDT Decap, MT - MDT Threshold Crossed, MH - MDT interface handle
             CD - Conditional Decap, MPLS - MPLS Decap, MF - MPLS Encap, EX - Extranet
             MoFE - MoFRR Enabled, MoFS - MoFRR State, MoFP - MoFRR Primary
             MoFB - MoFRR Backup, RPFID - RPF ID Set
Interface flags: F - Forward, A - Accept, IC - Internal Copy,
                NS - Negate Signal, DP - Don't Preserve, SP - Signal Present,
                II - Internal Interest, ID - Internal Disinterest, LI - Local Interest,
                LD - Local Disinterest, DI - Decapsulation Interface
                EI - Encapsulation Interface, MI - MDT Interface, LVIF - MPLS Encap,
                EX - Extranet, A2 - Secondary Accept, MT - MDT Threshold Crossed,
                MA - Data MDT Assigned, LMI - mLDP MDT Interface, TMI - P2MP-TE MDT Interface
                IRMI - IR MDT Interface

(*,224.0.0.0/4) RPF nbr: 10.0.1.1 Flags: L C P
  Up: 4w0d
  Outgoing Interface List
    Decapstunnel0 Flags: NS DI, Up: 4w0d

(*,224.0.0.0/24) Flags: D P
  Up: 4w0d

(*,224.0.1.39) Flags: S P
  Up: 4w0d

(*,224.0.1.40) Flags: S P
  Up: 4w0d
  Outgoing Interface List
    Loopback0 Flags: II LI, Up: 4w0d

(*,224.1.3.1) RPF nbr: 10.0.1.1 Flags: C
  Up: 20:31:54
  Incoming Interface List
    Decapstunnel0 Flags: A, Up: 20:31:54
  Outgoing Interface List
    GigabitEthernet0/0/0/3 Flags: F NS, Up: 20:31:54

(*,224.1.4.1) RPF nbr: 10.0.1.1 Flags: C
  Up: 20:31:16
  Incoming Interface List
    Decapstunnel0 Flags: A, Up: 20:31:16
  Outgoing Interface List
    GigabitEthernet0/0/0/5 Flags: F NS, Up: 20:31:16

(192.168.104.41,224.1.4.1) RPF nbr: 10.0.1.1 Flags: C
  Up: 00:02:45
  Incoming Interface List
    Decapstunnel0 Flags: A, Up: 00:02:45

(*,232.0.0.0/8) Flags: D P
  Up: 4w0d
```

Step 3 What are the incoming interface and OIL for (*,G) on the CE, PE, and P1 routers? G is the group for your pod (224.1.x.1 or 224.1.y.1).

Router	Incoming Interface	OIL
CE		
PE		
P1		

Step 4 Why are there no incoming interfaces on the P1 router for the (*,G) entry?

Verification

No verification is needed for this task.

Task 3: Shared Tree Formation—Sources

In this task, you will trigger some multicast traffic. Then you will observe the multicast routing tables on routers.

Activity Procedure

Complete these steps:

- Step 1** On the CE router add network 192.168.10x.0/24 or 192.168.10y.0/24 into the BGP routing process.
- Step 2** On the PE router remove the maximum prefix configuration configured in the previous exercise.
- Step 3** On the PE router, clear BGP sessions.
- The PE router (Cisco IOS XR) output should be similar to the following, taken from Pod 3:

```

RP/0/RSP0/CPU0:PE3# clear ip bgp *
RP/0/RSP0/CPU0:Jun 20 07:52:17.001 : bgp[1049]: %ROUTING-BGP-5-ADJCHANGE : neighbor
10.0.1.1 Down - User clear requested (CEASE notification sent - administrative reset)
(VRF: default) (AS: 64500)
RP/0/RSP0/CPU0:Jun 20 07:52:17.001 : bgp[1049]: %ROUTING-BGP-5-ADJCHANGE : neighbor
10.0.2.1 Down - User clear requested (CEASE notification sent - administrative reset)
(VRF: default) (AS: 64500)
RP/0/RSP0/CPU0:Jun 20 07:52:17.001 : bgp[1049]: %ROUTING-BGP-5-ADJCHANGE : neighbor
192.168.103.31 Down - User clear requested (CEASE notification sent - administrative
reset) (VRF: default) (AS: 64503)
RP/0/RSP0/CPU0:Jun 20 07:52:17.002 : bgp[1049]: %ROUTING-BGP-5-ADJCHANGE : neighbor
2001::db8:10:0:1::1 Down - User clear requested (CEASE notification sent -
administrative reset) (VRF: default) (AS: 64500)
RP/0/RSP0/CPU0:Jun 20 07:52:17.002 : bgp[1049]: %ROUTING-BGP-5-ADJCHANGE : neighbor
2001::db8:192:168:103::31 Down - User clear requested (CEASE notification sent -
administrative reset) (VRF: default) (AS: 64503)
RP/0/RSP0/CPU0:Jun 20 07:52:17.808 : bgp[1049]: %ROUTING-BGP-5-ADJCHANGE : neighbor
2001::db8:192:168:103::31 Up (VRF: default) (AS: 64503)
RP/0/RSP0/CPU0:Jun 20 07:52:17.919 : bgp[1049]: %ROUTING-BGP-5-ADJCHANGE : neighbor
192.168.103.31 Up (VRF: default) (AS: 64503)
RP/0/RSP0/CPU0:Jun 20 07:52:34.979 : bgp[1049]: %ROUTING-BGP-5-ADJCHANGE : neighbor
10.0.1.1 Up (VRF: default) (AS: 64500)
RP/0/RSP0/CPU0:Jun 20 07:52:39.054 : bgp[1049]: %ROUTING-BGP-5-ADJCHANGE : neighbor
10.0.2.1 Up (VRF: default) (AS: 64500)
RP/0/RSP0/CPU0:Jun 20 07:52:43.961 : bgp[1049]: %ROUTING-BGP-5-ADJCHANGE : neighbor
2001::db8:10:0:1::1 Up (VRF: default) (AS: 64500)

```

The PE router (Cisco IOS XE) output should be similar to the following, taken from Pod 4:

```

PE4# clear ip bgp *
*Jun 20 15:37:56.605: %BGP-5-ADJCHANGE: neighbor 10.0.1.1 Down User reset
*Jun 20 15:37:56.605: %BGP_SESSION-5-ADJCHANGE: neighbor 10.0.1.1 IPv4 Unicast topology
base removed from session User reset
*Jun 20 15:37:56.605: %BGP-5-ADJCHANGE: neighbor 10.0.2.1 Down User reset
*Jun 20 15:37:56.605: %BGP_SESSION-5-ADJCHANGE: neighbor 10.0.2.1 IPv4 Unicast topology
base removed from session User reset
*Jun 20 15:37:56.605: %BGP-5-ADJCHANGE: neighbor 2001::DB8:10:0:1::1 Down User reset
*Jun 20 15:37:56.605: %BGP_SESSION-5-ADJCHANGE: neighbor 2001::DB8:10:0:1::1 IPv6
Unicast topology base removed from session User reset
*Jun 20 15:37:56.605: %BGP-5-ADJCHANGE: neighbor 2001::DB8:192:168:104::41 Down User
reset
*Jun 20 15:37:56.605: %BGP_SESSION-5-ADJCHANGE: neighbor 2001::DB8:192:168:104::41 IPv6
Unicast topology base removed from session User reset
*Jun 20 15:37:56.605: %BGP-5-ADJCHANGE: neighbor 192.168.104.41 Down User reset
*Jun 20 15:37:56.605: %BGP_SESSION-5-ADJCHANGE: neighbor 192.168.104.41 IPv4 Unicast
topology base removed from session User reset
*Jun 20 15:38:07.870: %BGP-5-ADJCHANGE: neighbor 10.0.2.1 Up
*Jun 20 15:38:12.991: %BGP-5-ADJCHANGE: neighbor 2001::DB8:10:0:1::1 Up
*Jun 20 15:38:14.015: %BGP-5-ADJCHANGE: neighbor 10.0.1.1 Up

```

Step 4 Use Telnet to connect to the other pod CE router. Start the multicast traffic by pinging to your multicast group address. Use the GigabitEthernet 0/0 interface as a source interface. Send a large number of ICMP packets—100, for example. You should see that both receivers reply to the ping.

Students in Pod 3 will start the ping from the CE4 router. The CE4 router (Cisco IOS) output should be similar to the following, taken from Pod 3.

```

CE4# ping 224.1.3.1 repeat 100 source GigabitEthernet 0/0
Type escape sequence to abort.
Sending 100, 100-byte ICMP Echos to 224.1.3.1, timeout is 2 seconds:
Packet sent with a source address of 192.168.104.41

Reply to request 0 from 10.3.10.1, 1 ms
Reply to request 0 from 10.3.1.1, 1 ms
Reply to request 1 from 10.3.10.1, 1 ms
Reply to request 1 from 10.3.1.1, 1 ms
Reply to request 2 from 10.3.10.1, 16 ms
Reply to request 2 from 10.3.1.1, 16 ms
Reply to request 3 from 10.3.10.1, 1 ms
Reply to request 3 from 10.3.1.1, 1 ms
<... output omitted ...>

```

Students in Pod 4 will start the ping from the CE3 router. The CE3 router (Cisco IOS) output should be similar to the following, taken from Pod 4.

```

CE3# ping 224.1.4.1 repeat 100 source GigabitEthernet 0/0
Type escape sequence to abort.
Sending 100, 100-byte ICMP Echos to 224.1.4.1, timeout is 2 seconds:
Packet sent with a source address of 192.168.103.31

Reply to request 1 from 10.4.1.1, 1 ms
Reply to request 1 from 10.4.10.1, 1 ms
Reply to request 2 from 10.4.1.1, 1 ms
Reply to request 2 from 10.4.10.1, 1 ms
Reply to request 3 from 10.4.1.1, 1 ms
Reply to request 3 from 10.4.10.1, 1 ms
Reply to request 4 from 10.4.1.1, 1 ms
Reply to request 4 from 10.4.10.1, 1 ms
Reply to request 5 from 10.4.1.1, 1 ms
Reply to request 5 from 10.4.10.1, 1 ms
Reply to request 6 from 10.4.1.1, 1 ms
Reply to request 6 from 10.4.10.1, 1 ms
Reply to request 7 from 10.4.1.1, 1 ms

```

Step 5 Examine the multicast routing table on the CE router.

The CE router (Cisco IOS) output should be similar to the following, taken from Pod 3:

```

CE3# show ip mroute
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector
Outgoing interface flags: H - Hardware switched, A - Assert winner
Timers: Uptime/Expires
Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 224.1.4.1), 00:00:24/stopped, RP 10.0.1.1, flags: SPF
  Incoming interface: GigabitEthernet0/0, RPF nbr 192.168.103.30
  Outgoing interface list: Null

(192.168.103.31, 224.1.4.1), 00:00:24/00:02:53, flags: PFT
  Incoming interface: GigabitEthernet0/0, RPF nbr 0.0.0.0
  Outgoing interface list: Null

(*, 224.1.3.1), 00:00:47/00:02:59, RP 10.0.1.1, flags: SCL
  Incoming interface: GigabitEthernet0/0, RPF nbr 192.168.103.30
  Outgoing interface list:
    Loopback0, Forward/Sparse, 00:00:47/00:02:59

(*, 224.0.1.40), 00:00:47/00:02:57, RP 10.0.1.1, flags: SCL
  Incoming interface: GigabitEthernet0/0, RPF nbr 192.168.103.30
  Outgoing interface list:
    Loopback0, Forward/Sparse, 00:00:47/00:02:57

```

The CE router (Cisco IOS) output should be similar to the following, taken from Pod 4:

```

CE4# show ip mroute
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector
Outgoing interface flags: H - Hardware switched, A - Assert winner
Timers: Uptime/Expires
Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 224.1.1.4.1), 00:01:09/00:02:09, RP 10.0.1.1, flags: SCL
  Incoming interface: GigabitEthernet0/0, RPF nbr 192.168.104.40
  Outgoing interface list:
    Loopback0, Forward/Sparse, 00:01:09/00:02:09

(*, 224.1.1.3.1), 00:00:31/stopped, RP 10.0.1.1, flags: SPF
  Incoming interface: GigabitEthernet0/0, RPF nbr 192.168.104.40
  Outgoing interface list: Null

(192.168.104.41, 224.1.1.3.1), 00:00:31/00:02:46, flags: PFT
  Incoming interface: GigabitEthernet0/0, RPF nbr 0.0.0.0
  Outgoing interface list: Null

(*, 224.0.1.40), 00:01:09/00:02:41, RP 10.0.1.1, flags: SCL
  Incoming interface: GigabitEthernet0/0, RPF nbr 192.168.104.40
  Outgoing interface list:
    Loopback0, Forward/Sparse, 00:01:09/00:02:06

```

Step 6 Did the entry for your group change?

Step 7 Are any (S,G) entries present for your group? Why or why not?

Step 8 Use Telnet to connect the other pod PE router. Examine the multicast routing table.

Students in Pod 3 will connect to PE4 router. The PE4 router (Cisco IOS XE) output should be similar to the following, taken from Pod 3.

```

PE4# show ip mroute
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
L - Local, P - Pruned, R - RP-bit set, F - Register flag,
T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
U - URD, I - Received Source Specific Host Report,
Z - Multicast Tunnel, z - MDT-data group sender,
Y - Joined MDT-data group, y - Sending to MDT-data group,
G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
Q - Received BGP S-A Route, q - Sent BGP S-A Route,
V - RD & Vector, v - Vector, p - PIM Joins on route
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
Timers: Uptime/Expires
Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 224.1.4.1), 21:22:29/00:02:40, RP 10.0.1.1, flags: SCL
  Incoming interface: GigabitEthernet0/0/2, RPF nbr 192.168.41.1
  Outgoing interface list:
    GigabitEthernet0/0/0, Forward/Sparse, 00:16:35/00:03:09
    Loopback0, Forward/Sparse, 21:22:29/00:02:40

(*, 224.1.3.1), 00:15:01/stopped, RP 10.0.1.1, flags: SP
  Incoming interface: GigabitEthernet0/0/2, RPF nbr 192.168.41.1
  Outgoing interface list: Null

(192.168.104.41, 224.1.3.1), 00:02:33/00:00:26, flags: T
  Incoming interface: GigabitEthernet0/0/0, RPF nbr 192.168.104.41
  Outgoing interface list:
    GigabitEthernet0/0/2, Forward/Sparse, 00:00:28/00:03:01

(*, 224.0.1.40), 1w0d/00:03:09, RP 10.0.1.1, flags: SCL
  Incoming interface: GigabitEthernet0/0/2, RPF nbr 192.168.41.1
  Outgoing interface list:
    GigabitEthernet0/0/0, Forward/Sparse, 00:21:53/00:03:09

```

Students in Pod 4 will connect to the PE3 router. The PE3 router (Cisco IOS XR) output should be similar to the following, taken from Pod 4.

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

```
Fri Jun 20 09:06:04.091 UTC
```

```
IP Multicast Routing Information Base
```

```
Entry flags: L - Domain-Local Source, E - External Source to the Domain,
```

```
C - Directly-Connected Check, S - Signal, IA - Inherit Accept,
```

```
IF - Inherit From, D - Drop, ME - MDT Encap, EID - Encap ID,
```

```
MD - MDT Decap, MT - MDT Threshold Crossed, MH - MDT interface handle
```

```
CD - Conditional Decap, MPLS - MPLS Decap, MF - MPLS Encap, EX - Extranet
```

```
MoFE - MoFRR Enabled, MoFS - MoFRR State, MoFP - MoFRR Primary
```

```
MoFB - MoFRR Backup, RPFID - RPF ID Set
```

```
Interface flags: F - Forward, A - Accept, IC - Internal Copy,
```

```
NS - Negate Signal, DP - Don't Preserve, SP - Signal Present,
```

```
II - Internal Interest, ID - Internal Disinterest, LI - Local Interest,
```

```
LD - Local Disinterest, DI - Decapsulation Interface
```

```
EI - Encapsulation Interface, MI - MDT Interface, LVIF - MPLS Encap,
```

```
EX - Extranet, A2 - Secondary Accept, MT - MDT Threshold Crossed,
```

```
MA - Data MDT Assigned, LMI - mLDP MDT Interface, TMI - P2MP-TE MDT Interface
```

```
IRMI - IR MDT Interface
```

```
(* ,224.0.0.0/4) RPF nbr: 192.168.31.1 Flags: C P  
Up: 23:31:08
```

```
(* ,224.0.0.0/24) Flags: D P  
Up: 1w1d
```

```
(* ,224.0.1.39) Flags: S P  
Up: 1w1d
```

```
(* ,224.0.1.40) Flags: S P  
Up: 1w1d  
Outgoing Interface List  
GigabitEthernet0/0/0/0 Flags: II LI, Up: 1w1d
```

```
(* ,224.1.3.1) RPF nbr: 192.168.31.1 Flags: C  
Up: 22:18:46  
Incoming Interface List  
GigabitEthernet0/0/0/2 Flags: A, Up: 22:18:46  
Outgoing Interface List  
Loopback0 Flags: F IC NS II LI, Up: 22:18:46  
GigabitEthernet0/0/0/0 Flags: F NS, Up: 01:13:11
```

```
(192.168.103.31,224.1.4.1) RPF nbr: 192.168.103.31 Flags:  
Up: 00:00:28  
Incoming Interface List  
GigabitEthernet0/0/0/0 Flags: A, Up: 00:00:28  
Outgoing Interface List  
GigabitEthernet0/0/0/2 Flags: F NS, Up: 00:00:28
```

```
(* ,232.0.0.0/8) Flags: D P  
Up: 1w1d
```

Step 9 What is the incoming interface for the (S,G) entry for your group?

Step 10 Which interfaces are present in the OIL for the (S,G) entry for your group?

Step 11 Why is the OIL of the (*,G) entry for your group empty?

Step 12 Interrupt the multicast ping from the other pod CE router. This action is usually done using the Shift-Ctrl-6 key sequence twice.

Students in Pod 3 will start a ping from the CE4 router. The CE4 router (Cisco IOS) output should be similar to the following, taken from Pod 3.

```
<... output omitted ...>
Reply to request 6 from 10.3.10.1, 1 ms
Reply to request 6 from 10.3.1.1, 1 ms
Reply to request 7 from 10.3.10.1, 1 ms
Reply to request 7 from 10.3.1.1, 1 ms
Reply to request 8 from 10.3.10.1, 1 ms
Reply to request 8 from 10.3.1.1, 1 ms
Reply to request 9 from 10.3.10.1, 1 ms
Reply to request 9 from 10.3.1.1, 1 ms <Shift + Ctrl + 6, Shift + Ctrl + 6>
CE4#
```

Students in Pod 4 will start a ping from the CE3 router. The CE3 router (Cisco IOS) output should be similar to the following, taken from Pod 4.

```
<... output omitted ...>
Reply to request 30 from 10.4.1.1, 1 ms
Reply to request 30 from 10.4.10.1, 1 ms
Reply to request 31 from 10.4.1.1, 1 ms
Reply to request 31 from 10.4.10.1, 1 ms
Reply to request 32 from 10.4.1.1, 1 ms
Reply to request 32 from 10.4.10.1, 1 ms
Reply to request 33 from 10.4.1.1, 1 ms
Reply to request 33 from 10.4.10.1, 1 ms
Reply to request 34 from 10.4.1.1, 1 ms
Reply to request 34 from 10.4.10.1, 1 ms
Reply to request 35 from 10.4.1.1, 1 ms
Reply to request 35 from 10.4.10.1, 1 ms <Shift + Ctrl + 6, Shift + Ctrl + 6>
CE3#
```

Verification

No verification is needed for this task.

Task 4: Switching to the SPT

In this task, you will configure the last-hop routers to switch to SPT immediately after the first packet is received over the shared tree.

Activity Procedure

Complete these steps:

- Step 1** Configure the CE and PE routers to switch to SPT immediately after the first packet arrives over the shared tree.
- Step 2** Use Telnet to connect to the other pod CE router. Start the multicast traffic by pinging to your multicast group address. Use the GigabitEthernet0/0 interface as a source interface. Send a large number of ICMP packets—100, for example. You should see that both receivers reply to the ping.

Students in Pod 3 will start a ping from the CE4 router. The CE4 router (Cisco IOS) output should be similar to the following, taken from Pod 3.

```
CE4# ping 224.1.3.1 repeat 100 source GigabitEthernet 0/0
Type escape sequence to abort.
Sending 100, 100-byte ICMP Echos to 224.1.3.1, timeout is 2 seconds:
Packet sent with a source address of 192.168.104.41

Reply to request 0 from 10.3.10.1, 1 ms
Reply to request 0 from 10.3.1.1, 1 ms
Reply to request 1 from 10.3.10.1, 1 ms
Reply to request 1 from 10.3.1.1, 1 ms
Reply to request 2 from 10.3.10.1, 16 ms
Reply to request 2 from 10.3.1.1, 16 ms
Reply to request 3 from 10.3.10.1, 1 ms
Reply to request 3 from 10.3.1.1, 1 ms
<... output omitted ...>
```

Students in Pod 4 will start a ping from the CE3 router. The CE3 router (Cisco IOS) output should be similar to the following, taken from Pod 4.

```
CE3# ping 224.1.4.1 repeat 100 source GigabitEthernet 0/0
Type escape sequence to abort.
Sending 100, 100-byte ICMP Echos to 224.1.4.1, timeout is 2 seconds:
Packet sent with a source address of 192.168.103.31

Reply to request 1 from 10.4.1.1, 1 ms
Reply to request 1 from 10.4.10.1, 1 ms
Reply to request 2 from 10.4.1.1, 1 ms
Reply to request 2 from 10.4.10.1, 1 ms
Reply to request 3 from 10.4.1.1, 1 ms
Reply to request 3 from 10.4.10.1, 1 ms
Reply to request 4 from 10.4.1.1, 1 ms
Reply to request 4 from 10.4.10.1, 1 ms
Reply to request 5 from 10.4.1.1, 1 ms
Reply to request 5 from 10.4.10.1, 1 ms
Reply to request 6 from 10.4.1.1, 1 ms
Reply to request 6 from 10.4.10.1, 1 ms
Reply to request 7 from 10.4.1.1, 1 ms
```

- Step 3** Examine the multicast routing table on the CE router.

The CE router (Cisco IOS) output should be similar to the following, taken from Pod 3:

```

CE3# show ip mroute
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector
Outgoing interface flags: H - Hardware switched, A - Assert winner
Timers: Uptime/Expires
Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 224.1.4.1), 00:00:28/stopped, RP 10.0.1.1, flags: SPF
  Incoming interface: GigabitEthernet0/0, RPF nbr 192.168.103.30
  Outgoing interface list: Null

(192.168.103.31, 224.1.4.1), 00:00:28/00:02:49, flags: PFT
  Incoming interface: GigabitEthernet0/0, RPF nbr 0.0.0.0
  Outgoing interface list: Null

(*, 224.1.3.1), 00:44:05/stopped, RP 10.0.1.1, flags: SJCL
  Incoming interface: GigabitEthernet0/0, RPF nbr 192.168.103.30
  Outgoing interface list:
    Loopback0, Forward/Sparse, 00:44:05/00:02:45

(192.168.104.41, 224.1.3.1), 00:00:24/00:02:35, flags: LJT
  Incoming interface: GigabitEthernet0/0, RPF nbr 192.168.103.30
  Outgoing interface list:
    Loopback0, Forward/Sparse, 00:00:24/00:02:45

(*, 224.0.1.40), 00:44:05/00:02:47, RP 10.0.1.1, flags: SJCL
  Incoming interface: GigabitEthernet0/0, RPF nbr 192.168.103.30
  Outgoing interface list:
    Loopback0, Forward/Sparse, 00:44:05/00:02:47

```

The CE router (Cisco IOS) output should be similar to the following, taken from Pod 4:

```

CE4# show ip mroute
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector
Outgoing interface flags: H - Hardware switched, A - Assert winner
Timers: Uptime/Expires
Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 224.1.4.1), 00:44:21/stopped, RP 10.0.1.1, flags: SJCL
  Incoming interface: GigabitEthernet0/0, RPF nbr 192.168.104.40
  Outgoing interface list:
    Loopback0, Forward/Sparse, 00:44:21/00:02:08

(192.168.103.31, 224.1.4.1), 00:00:33/00:02:26, flags: LJT
  Incoming interface: GigabitEthernet0/0, RPF nbr 192.168.104.40
  Outgoing interface list:
    Loopback0, Forward/Sparse, 00:00:33/00:02:26

(*, 224.1.3.1), 00:00:34/stopped, RP 10.0.1.1, flags: SPF
  Incoming interface: GigabitEthernet0/0, RPF nbr 192.168.104.40
  Outgoing interface list: Null

(192.168.104.41, 224.1.3.1), 00:00:34/00:02:43, flags: PFT
  Incoming interface: GigabitEthernet0/0, RPF nbr 0.0.0.0
  Outgoing interface list: Null

(*, 224.0.1.40), 00:44:21/00:02:57, RP 10.0.1.1, flags: SJCL
  Incoming interface: GigabitEthernet0/0, RPF nbr 192.168.104.40
  Outgoing interface list:
    Loopback0, Forward/Sparse, 00:44:21/00:02:04

```

Step 4 Are any (S,G) entries present for your group? Why or why not?

Step 5 Examine the multicast routing table on the PE router.

The PE router (Cisco IOS XR) output should be similar to the following, taken from Pod 3:

```
RP/0/RSP0/CPU0:PE3# show mrib route
Fri Jun 20 08:57:20.885 UTC
```

IP Multicast Routing Information Base

Entry flags: L - Domain-Local Source, E - External Source to the Domain,
C - Directly-Connected Check, S - Signal, IA - Inherit Accept,
IF - Inherit From, D - Drop, ME - MDT Encap, EID - Encap ID,
MD - MDT Decap, MT - MDT Threshold Crossed, MH - MDT interface handle
CD - Conditional Decap, MPLS - MPLS Decap, MF - MPLS Encap, EX - Extranet
MoFE - MoFRR Enabled, MoFS - MoFRR State, MoFP - MoFRR Primary
MoFB - MoFRR Backup, RPFID - RPF ID Set

Interface flags: F - Forward, A - Accept, IC - Internal Copy,
NS - Negate Signal, DP - Don't Preserve, SP - Signal Present,
II - Internal Interest, ID - Internal Disinterest, LI - Local Interest,
LD - Local Disinterest, DI - Decapsulation Interface
EI - Encapsulation Interface, MI - MDT Interface, LVIF - MPLS Encap,
EX - Extranet, A2 - Secondary Accept, MT - MDT Threshold Crossed,
MA - Data MDT Assigned, LMI - mLDP MDT Interface, TMI - P2MP-TE MDT Interface
IRMI - IR MDT Interface

```
{*,224.0.0.0/4) RPF nbr: 192.168.31.1 Flags: C P
Up: 23:22:25
```

```
{*,224.0.0.0/24) Flags: D P
Up: 1w1d
```

```
{*,224.0.1.39) Flags: S P
Up: 1w1d
```

```
{*,224.0.1.40) Flags: S P
Up: 1w1d
Outgoing Interface List
GigabitEthernet0/0/0/0 Flags: II LI, Up: 1w1d
```

```
{*,224.1.3.1) RPF nbr: 192.168.31.1 Flags: C
Up: 22:10:03
Incoming Interface List
GigabitEthernet0/0/0/2 Flags: A NS, Up: 22:10:03
Outgoing Interface List
Loopback0 Flags: F IC NS II LI, Up: 22:10:03
GigabitEthernet0/0/0/0 Flags: F NS, Up: 01:04:28
```

```
{192.168.104.41,224.1.3.1) RPF nbr: 192.168.134.40 Flags:
Up: 00:09:22
Incoming Interface List
GigabitEthernet0/0/0/1 Flags: A, Up: 00:09:22
Outgoing Interface List
Loopback0 Flags: F IC NS, Up: 00:09:22
GigabitEthernet0/0/0/0 Flags: F NS, Up: 00:09:22
```

```
{192.168.103.31,224.1.4.1) RPF nbr: 192.168.103.31 Flags:
Up: 00:09:26
Incoming Interface List
GigabitEthernet0/0/0/0 Flags: A, Up: 00:09:26
Outgoing Interface List
GigabitEthernet0/0/0/1 Flags: F NS, Up: 00:09:22
```

```
{*,232.0.0.0/8) Flags: D P
Up: 1w1d
```

The PE router (Cisco IOS XE) output should be similar to the following, taken from Pod 4:

```

PE4# show ip mroute
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
Timers: Uptime/Expires
Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 224.1.4.1), 22:04:34/00:03:20, RP 10.0.1.1, flags: SJCL
  Incoming interface: GigabitEthernet0/0/2, RPF nbr 192.168.41.1
  Outgoing interface list:
    GigabitEthernet0/0/0, Forward/Sparse, 00:58:40/00:03:20
    Loopback0, Forward/Sparse, 22:04:34/00:02:40

(192.168.103.31, 224.1.4.1), 00:00:17/00:02:42, flags: LT
  Incoming interface: GigabitEthernet0/0/1, RPF nbr 192.168.134.30
  Outgoing interface list:
    Loopback0, Forward/Sparse, 00:00:17/00:02:42
    GigabitEthernet0/0/0, Forward/Sparse, 00:00:17/00:03:20

(*, 224.1.3.1), 00:04:33/stopped, RP 10.0.1.1, flags: SP
  Incoming interface: GigabitEthernet0/0/2, RPF nbr 192.168.41.1
  Outgoing interface list: Null

(192.168.104.41, 224.1.3.1), 00:01:29/00:01:30, flags: T
  Incoming interface: GigabitEthernet0/0/0, RPF nbr 192.168.104.41
  Outgoing interface list:
    GigabitEthernet0/0/1, Forward/Sparse, 00:01:29/00:03:00

(*, 224.0.1.40), 1w0d/00:03:23, RP 10.0.1.1, flags: SJCL
  Incoming interface: GigabitEthernet0/0/2, RPF nbr 192.168.41.1
  Outgoing interface list:
    GigabitEthernet0/0/0, Forward/Sparse, 01:03:59/00:03:23

```

Step 6 Are any (S,G) entries present for your group? Why or why not?

Step 7 Which interface is used as the incoming interface for the (S,G) entry for your group?

Step 8 Remove the simulated multicast receivers for group 224.1.x.1 (or 224.1.y.0) from the Loopback0 interface on the CE and PE routers.

Verification

No verification is needed for this task.

Hardware Lab 6: Implement PIM-SM Enhancements

Activity Objective

In this lab activity, you will configure and monitor PIM-SSM on the CE and PE routers.


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

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

- Implement and monitor PIM-SSM

Visual Objective

The figure illustrates what you will accomplish in this activity.



Command List

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

Cisco IOS/IOS XE Software Commands

Command	Description
configure terminal	Enters configuration mode.
interface interface	Enters interface configuration mode.
ip access-list standard acl_name	Creates a standard ACL and enters access list configuration mode.
ip igmp join-group group_address [source source]	Configures an interface on the router to join the specified group or channel.
ip pim bidir-enable	Globally enables BIDIR-PIM.
ip pim rp-address RP_address [bidir]	Statically configures the address of a PIM RP for multicast groups and enables BIDIR-PIM.
ip pim ssm range acl_name	Enables PIM-SSM for specified groups.

Command	Description
[permit deny] host <i>IP_address</i>	Creates a standard ACL entry under access list configuration mode.
ping <i>dest_IP</i> <i>source</i> <i>source_IP</i>	Verifies connectivity between the source IP and the destination IP (IPv4 and IPv6).
show ip mroute	Displays the contents of the multicast route (mroute) table.
show ip pim interface	Displays information about interfaces that are configured for PIM.
show ip pim interface <i>df</i>	Displays the IP address of the elected DF for each RP of an interface.
show ip pim neighbor	Displays information about PIM neighbors discovered by PIMv1 router query messages or PIMv2 hello messages.

Cisco IOS XR Software Commands

Command	Description
address-family ipv4	Enters an IPv4 address family under a specific configuration mode.
commit	Commits changes to the running configuration.
configure	Enters configuration mode.
interface <i>interface</i>	Enters interface configuration mode.
ipv4 access-list <i>acl_name</i>	Creates an ACL and enters access list configuration mode.
join-group <i>group_address</i> [<i>source</i>]	Configures an interface on the router to join the specified group or channel.
multicast-routing	Enters multicast routing configuration mode.
[permit deny] host <i>IP_address</i>	Creates a standard ACL entry under access list configuration mode.
ping <i>dest_IP</i> <i>source</i> <i>source_IP</i>	Verifies connectivity between the source IP and the destination IP (IPv4 and IPv6).
router igmp	Enters IGMP configuration mode.
router pim	Enters PIM configuration mode.
rp-address <i>RP_address</i> [<i>bidir</i>]	Statically configures the address of a PIM RP for multicast groups and enables BIDIR-PIM.

Command	Description
show mrib route	Displays the contents of the multicast routing (mroute) table.
show pim df winner	Displays the IP address of the elected DF for each RP of an interface.
show pim interface	Displays information about interfaces that are configured for PIM.
show pim neighbor	Displays information about PIM neighbors discovered by PIMv1 router query messages or PIMv2 hello messages.
ssm range <i>acl_name</i>	Enables PIM-SSM for specified groups under multicast configuration mode.

Task 1: Implement PIM-SSM

In this task, you will enable PIM-SSM on the CE and PE routers in your pod. Multicast routing and PIM should be already enabled from the previous lab exercise.

Activity Procedure

Complete these steps:

- Step 1** Enable PIM-SSM on the PE and CE router. Configure the SSM address range to include the 224.1.x.1 and 224.1.y.1 multicast groups.

Note Both pods from the same team should be finished with the previous step at this point.

- Step 2** Verify the PIM state on interfaces on the PE router.

The PE router (Cisco IOS XR) output should be similar to the following, taken from Pod 3:

```
RP/0/RSP0/CP00:PE3# show pim interface
PIM interfaces in VRF default
Address                Interface                PIM  Nbr   Hello  DR    DR
                        Count Intvl  Prior
10.3.1.1                Loopback0                on   1     30     1     this system
192.168.103.30          GigabitEthernet0/0/0/0   on   2     30     1
192.168.103.31
192.168.134.30          GigabitEthernet0/0/0/1   on   2     30     1
192.168.134.40
192.168.31.30           GigabitEthernet0/0/0/2   on   2     30     1     this system
192.168.32.30           GigabitEthernet0/0/0/3   on   2     30     1     this system
```

The PE router (Cisco IOS XE) output should be similar to the following, taken from Pod 4:

```
PE4# show ip pim interface
```

Address	Interface	Ver/ Mode	Nbr Count	Query Intvl	DR Prior	DR
192.168.104.40	GigabitEthernet0/0/0	v2/S	1	30	1	192.168.104.41
10.4.1.1	Loopback0	v2/S	0	30	1	10.4.1.1
192.168.134.40	GigabitEthernet0/0/1	v2/S	1	30	1	192.168.134.40
192.168.41.40	GigabitEthernet0/0/2	v2/S	1	30	1	192.168.41.40
192.168.42.40	GigabitEthernet0/0/3	v2/S	1	30	1	192.168.42.40

Step 3 Verify the PIM neighbors on the PE router.

The PE router (Cisco IOS XR) output should be similar to the following, taken from Pod 3:

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

PIM neighbors in VRF default

Flag: B - Bidir capable, P - Proxy capable, DR - Designated Router,

E - ECMP Redirect capable

* indicates the neighbor created for this router

Neighbor Address	Interface	Uptime	Expires	DR pri	Flags
10.3.1.1*	Loopback0	1d03h	00:01:32	1 (DR)	B P E
192.168.103.30*	GigabitEthernet0/0/0/0	1w1d	00:01:37	1	B P E
192.168.103.31	GigabitEthernet0/0/0/0	1d02h	00:01:40	1 (DR)	P
192.168.134.30*	GigabitEthernet0/0/0/1	1d03h	00:01:38	1	B P E
192.168.134.40	GigabitEthernet0/0/0/1	1d02h	00:01:22	1 (DR)	P
192.168.31.1	GigabitEthernet0/0/0/2	1d03h	00:01:38	1	B P
192.168.31.30*	GigabitEthernet0/0/0/2	1d03h	00:01:23	1 (DR)	B P E
192.168.32.2	GigabitEthernet0/0/0/3	1d03h	00:01:25	1	B P
192.168.32.30*	GigabitEthernet0/0/0/3	1d03h	00:01:42	1 (DR)	B P E

The PE router (Cisco IOS XE) output should be similar to the following, taken from Pod 4:

```
PE4# show ip pim neighbor
```

PIM Neighbor Table

Mode: B - Bidir Capable, DR - Designated Router, N - Default DR Priority,

P - Proxy Capable, S - State Refresh Capable, G - GenID Capable

Neighbor Address	Interface	Uptime/Expires	Ver	DR Prio/Mode
192.168.104.41	GigabitEthernet0/0/0	1d02h/00:01:16	v2	1 / DR S P G
192.168.134.30	GigabitEthernet0/0/1	1d02h/00:01:15	v2	1 / P G
192.168.41.1	GigabitEthernet0/0/2	1d02h/00:01:27	v2	1 / P G
192.168.42.2	GigabitEthernet0/0/3	1d02h/00:01:43	v2	1 / P G

Step 4 Simulate multicast receivers for group 224.l.x.1 (or 224.l.y.1) on the Loopback0 interface on the CE and PE routers. The routers should be interested only in traffic coming from the other pod CE router GigabitEthernet0/0 interface.

NOTE: Throughout the lab exercise, use the **ip igmp join-group** command on Cisco IOS and IOS-XE Software under interface configuration mode to simulate multicast receivers. On the Cisco IOS XR Software, use the **join-group** command under interface configuration mode under router IGMP configuration mode.

Step 5 Examine the multicast routing table on the PE router.

The PE router (Cisco IOS XR) output should be similar to the following, taken from Pod 3:

```

RP/0/RSP0/CPU0:PE3# show mrib route
IP Multicast Routing Information Base
Entry flags: L - Domain-Local Source, E - External Source to the Domain,
             C - Directly-Connected Check, S - Signal, IA - Inherit Accept,
             IF - Inherit From, D - Drop, ME - MDT Encap, EID - Encap ID,
             MD - MDT Decap, MT - MDT Threshold Crossed, MH - MDT interface handle
             CD - Conditional Decap, MPLS - MPLS Decap, MF - MPLS Encap, EX - Extranet
             MoFE - MoFRR Enabled, MoFS - MoFRR State, MoFP - MoFRR Primary
             MoFB - MoFRR Backup, RPFID - RPF ID Set
Interface flags: F - Forward, A - Accept, IC - Internal Copy,
                NS - Negate Signal, DP - Don't Preserve, SP - Signal Present,
                II - Internal Interest, ID - Internal Disinterest, LI - Local Interest,
                LD - Local Disinterest, DI - Decapsulation Interface
                EI - Encapsulation Interface, MI - MDT Interface, LVIF - MPLS Encap,
                EX - Extranet, A2 - Secondary Accept, MT - MDT Threshold Crossed,
                MA - Data MDT Assigned, LMI - mLDP MDT Interface, TMI - P2MP-TE MDT Interface
                IRMI - IR MDT Interface

(*,224.0.0.0/4) RPF nbr: 192.168.31.1 Flags: C P
  Up: 1d02h

(*,224.0.0.0/24) Flags: D P
  Up: 1w1d

(*,224.0.1.39) Flags: S P
  Up: 1w1d

(*,224.0.1.40) Flags: S P
  Up: 1w1d
  Outgoing Interface List
    GigabitEthernet0/0/0/0 Flags: II LI, Up: 1w1d

(192.168.104.41,224.1.3.1) RPF nbr: 192.168.134.40 Flags:
  Up: 00:04:25
  Incoming Interface List
    GigabitEthernet0/0/0/1 Flags: A, Up: 00:04:25
  Outgoing Interface List
    Loopback0 Flags: F IC NS II LI, Up: 00:04:25
    GigabitEthernet0/0/0/0 Flags: F NS, Up: 00:00:57

(192.168.103.31,224.1.4.1) RPF nbr: 192.168.103.31 Flags:
  Up: 00:03:30
  Incoming Interface List
    GigabitEthernet0/0/0/0 Flags: A, Up: 00:03:30
  Outgoing Interface List
    GigabitEthernet0/0/0/1 Flags: F NS, Up: 00:03:30

```

The PE router (Cisco IOS XE) output should be similar to the following, taken from Pod 4:

```

PE4# show ip mroute
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
Timers: Uptime/Expires
Interface state: Interface, Next-Hop or VCD, State/Mode

(192.168.103.31, 224.1.4.1), 00:03:18/00:03:09, flags: sLTI
  Incoming interface: GigabitEthernet0/0/1, RPF nbr 192.168.134.30
  Outgoing interface list:
    Loopback0, Forward/Sparse, 00:02:25/00:00:34
    GigabitEthernet0/0/0, Forward/Sparse, 00:03:18/00:03:09

(192.168.104.41, 224.1.3.1), 00:04:13/00:03:11, flags: sT
  Incoming interface: GigabitEthernet0/0/0, RPF nbr 192.168.104.41
  Outgoing interface list:
    GigabitEthernet0/0/1, Forward/Sparse, 00:04:13/00:03:11

(*, 224.0.1.40), 1w0d/00:02:59, RP 10.0.1.1, flags: SJCL
  Incoming interface: GigabitEthernet0/0/2, RPF nbr 192.168.41.1
  Outgoing interface list:
    GigabitEthernet0/0/0, Forward/Sparse, 04:09:58/00:03:27

```

Step 6 What is the incoming interface for the (S,G) entry for your group?

Step 7 What is the OIL for the (S,G) entry for your group?

Step 8 Use Telnet to connect to the other pod PE router and examine the multicast routing table.
 Students in Pod 3 will examine multicast routing table from PE4 router. The PE4 router (Cisco IOS XE) output should be similar to the following, taken from Pod 3.

```

PE4# show ip mroute
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
L - Local, P - Pruned, R - RP-bit set, F - Register flag,
T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
U - URD, I - Received Source Specific Host Report,
Z - Multicast Tunnel, z - MDT-data group sender,
Y - Joined MDT-data group, y - Sending to MDT-data group,
G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
Q - Received BGP S-A Route, q - Sent BGP S-A Route,
V - RD & Vector, v - Vector, p - PIM Joins on route
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
Timers: Uptime/Expires
Interface state: Interface, Next-Hop or VCD, State/Mode

{192.168.103.31, 224.1.4.1}, 00:03:18/00:03:09, flags: sLTl
  Incoming interface: GigabitEthernet0/0/1, RPF nbr 192.168.134.30
  Outgoing interface list:
    Loopback0, Forward/Sparse, 00:02:25/00:00:34
    GigabitEthernet0/0/0, Forward/Sparse, 00:03:18/00:03:09

{192.168.104.41, 224.1.3.1}, 00:04:13/00:03:11, flags: sT
  Incoming interface: GigabitEthernet0/0/0, RPF nbr 192.168.104.41
  Outgoing interface list:
    GigabitEthernet0/0/1, Forward/Sparse, 00:04:13/00:03:11

(*, 224.0.1.40), 1w0d/00:02:59, RP 10.0.1.1, flags: SJCL
  Incoming interface: GigabitEthernet0/0/2, RPF nbr 192.168.41.1
  Outgoing interface list:
    GigabitEthernet0/0/0, Forward/Sparse, 04:09:58/00:03:27

```

Students in Pod 4 will examine multicast routing table from the PE3 router. The PE3 router (Cisco IOS XR) output should be similar to the following, taken from Pod 4.

```

RP/0/RSP0/CPU0:PE3# show mrib route
IP Multicast Routing Information Base
Entry flags: L - Domain-Local Source, E - External Source to the Domain,
             C - Directly-Connected Check, S - Signal, IA - Inherit Accept,
             IF - Inherit From, D - Drop, ME - MDT Encap, EID - Encap ID,
             MD - MDT Decap, MT - MDT Threshold Crossed, MH - MDT interface handle
             CD - Conditional Decap, MPLS - MPLS Decap, MF - MPLS Encap, EX - Extranet
             MoFE - MoFRR Enabled, MoFS - MoFRR State, MoFP - MoFRR Primary
             MoFB - MoFRR Backup, RPFID - RPF ID Set
Interface flags: F - Forward, A - Accept, IC - Internal Copy,
                NS - Negate Signal, DP - Don't Preserve, SP - Signal Present,
                II - Internal Interest, ID - Internal Disinterest, LI - Local Interest,
                LD - Local Disinterest, DI - Decapsulation Interface
                EI - Encapsulation Interface, MI - MDT Interface, LVIF - MPLS Encap,
                EX - Extranet, A2 - Secondary Accept, MT - MDT Threshold Crossed,
                MA - Data MDT Assigned, LMI - mLDP MDT Interface, TMI - P2MP-TE MDT Interface
                IRMI - IR MDT Interface

(*,224.0.0.0/4) RPF nbr: 192.168.31.1 Flags: C P
  Up: 1d02h

(*,224.0.0.0/24) Flags: D P
  Up: 1wld

(*,224.0.1.39) Flags: S P
  Up: 1wld

(*,224.0.1.40) Flags: S P
  Up: 1wld
  Outgoing Interface List
    GigabitEthernet0/0/0/0 Flags: II LI, Up: 1wld

(192.168.104.41,224.1.3.1) RPF nbr: 192.168.134.40 Flags:
  Up: 00:04:25
  Incoming Interface List
    GigabitEthernet0/0/0/1 Flags: A, Up: 00:04:25
  Outgoing Interface List
    Loopback0 Flags: F IC NS II LI, Up: 00:04:25
    GigabitEthernet0/0/0/0 Flags: F NS, Up: 00:00:57

(192.168.103.31,224.1.4.1) RPF nbr: 192.168.103.31 Flags:
  Up: 00:03:30
  Incoming Interface List
    GigabitEthernet0/0/0/0 Flags: A, Up: 00:03:30
  Outgoing Interface List
    GigabitEthernet0/0/0/1 Flags: F NS, Up: 00:03:30

```

You should see that routers created SPT across all routers between the source and receivers.

Step 9 What is the incoming interface for the (S,G) entry for your group?

Step 10 What is the OIL for the (S,G) entry for your group?

Step 11 Use Telnet to connect to the other pod CE router. Start the multicast traffic by pinging to your multicast group address. Use the GigabitEthernet0/0 interface as a source interface. Send a large number of ICMP packets—100, for example. You should see that both receivers reply to the ping.

Students in Pod 3 will start the ping from the CE4 router. The CE4 router (Cisco IOS) output should be similar to the following, taken from Pod 3.

```

CE4# ping 224.1.3.1 repeat 100 source GigabitEthernet 0/0
Type escape sequence to abort.
Sending 100, 100-byte ICMP Echos to 224.1.3.1, timeout is 2 seconds:
Packet sent with a source address of 192.168.104.41

Reply to request 0 from 10.3.10.1, 1 ms
Reply to request 0 from 10.3.1.1, 1 ms
Reply to request 1 from 10.3.10.1, 1 ms
Reply to request 1 from 10.3.1.1, 1 ms
Reply to request 2 from 10.3.10.1, 16 ms
Reply to request 2 from 10.3.1.1, 16 ms
Reply to request 3 from 10.3.10.1, 1 ms
Reply to request 3 from 10.3.1.1, 1 ms
<... output omitted ...>

```

Students in Pod 4 will start the ping from the CE3 router. The CE3 router (Cisco IOS) output should be similar to the following, taken from Pod 4.

```

CE3# ping 224.1.4.1 repeat 100 source GigabitEthernet 0/0
Type escape sequence to abort.
Sending 100, 100-byte ICMP Echos to 224.1.4.1, timeout is 2 seconds:
Packet sent with a source address of 192.168.103.31

Reply to request 1 from 10.4.1.1, 1 ms
Reply to request 1 from 10.4.10.1, 1 ms
Reply to request 2 from 10.4.1.1, 1 ms
Reply to request 2 from 10.4.10.1, 1 ms
Reply to request 3 from 10.4.1.1, 1 ms
Reply to request 3 from 10.4.10.1, 1 ms
Reply to request 4 from 10.4.1.1, 1 ms
Reply to request 4 from 10.4.10.1, 1 ms
Reply to request 5 from 10.4.1.1, 1 ms
Reply to request 5 from 10.4.10.1, 1 ms
Reply to request 6 from 10.4.1.1, 1 ms
Reply to request 6 from 10.4.10.1, 1 ms
Reply to request 7 from 10.4.1.1, 1 ms

```

Step 12 Remove the simulated multicast receivers for group 224.1.x.1 (or 224.1.y.0) from the Loopback0 interface on the CE and PE routers.

Step 13 Disable PIM-SSM on the PE and CE routers.

Verification

No verification is needed for this task.

Hardware Lab 7: Implement Rendezvous Point Distribution

Activity Objective

In this activity, you will use PIM-SM. Avoiding the configuration of static RP information, you will choose the standard bootstrap mechanism to verify the redundant setup of BSR routers and RPs.

You also will configure two RPs with the same IP address, sharing the same range of groups. This action will create the Anycast RP solution that will require a simple MSDP configuration.


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

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

- Enable OSPF Routing
- Enable BSR
- Enable Anycast RP

Visual Objective

The figure illustrates what you will accomplish in this activity.



Command List

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

Cisco IOS/IOS XE Software Commands

Command	Description
clear ip pim rp-mapping	Clears the group-to-RP mapping table.
configure terminal	Enters configuration mode.
interface interface	Enters interface configuration mode.
ip igmp join-group group_address [source source]	Configures an interface on the router to join the specified group or channel.
ip msdp peer IP-address connect-source interface	Enables an MSDP peer.
ip pim bsr-candidate interface	Enables a candidate BSR.
ip pim rp-address RP_address	Statically configures the address of a PIM RP for multicast groups.

Command	Description
ip pim rp-candidate <i>interface</i>	Enables a BSR candidate RP.
ip pim sparse-mode	Enables PIM-SM on the interface.
network <i>network wildcard-mask area area-ID</i>	Enables an OSPF on the interface.
neighbor <i>peer-address shutdown</i>	Disables BGP neighbor adjacency.
ping <i>dest_IP source source_IP</i>	Verifies connectivity between the source IP and the destination IP (IPv4 and IPv6).
router bgp <i>AS-number</i>	Enables the BGP routing protocol and enters router BGP routing protocol configuration mode.
router ospf <i>process-ID</i>	Enables the OSPF routing protocol and enters router OSPF routing protocol configuration mode.
show ip mroute	Displays the contents of the multicast routing (mroute) table.
show ip msdp peer	Displays MSDP peer adjacency information.
show ip pim rp mapping	Displays group-to-RP mapping table.

Cisco IOS XR Software Commands

Command	Description
address-family <i>ipv4</i>	Enters an IPv4 address family under a specific configuration mode.
bsr candidate-bsr <i>IP-address</i>	In router PIM configuration mode, enables a candidate BSR.
bsr candidate-rp <i>IP-address</i>	In router PIM configuration mode, enables a BSR candidate RP.
clear pim bsr	Clears the BSR group-to-RP mapping table.
commit	Commits changes to the running configuration.
configure	Enters configuration mode.
connect-source <i>interface</i>	Enables an MSDP source interface.
enable	In router PIM configuration mode, enables PIM-SM on the interface.
interface <i>interface</i>	Enters interface configuration mode.

Command	Description
join-group <i>group_address</i> [<i>source source</i>]	Configures an interface on the router to join the specified group or channel.
mcast-routing	Enters multicast routing configuration mode.
peer <i>IP-address</i>	Enables an MSDP peer IP address.
ping <i>dest_IP source source_IP</i>	Verifies connectivity between the source IP and the destination IP (IPv4 and IPv6).
router bgp <i>AS-number</i>	Enables the BGP routing protocol and enters router BGP routing protocol configuration mode.
router igmp	Enters IGMP configuration mode.
router ospf <i>process-ID</i>	Enables the OSPF routing protocol and enters router OSPF routing protocol configuration mode.
router msdp	Enters MSDP configuration mode.
router pim	Enters PIM configuration mode.
rp-address <i>RP_address</i>	Statically configures the address of a PIM RP for multicast groups.
show mrib route	Displays the contents of the multicast routing (mroute) table.
show msdp peer	Displays MSDP peer adjacency information.
show pim group-map	Displays the group-to-RP mapping table.
shutdown	Disables an interface.

Task 1: Enable OSPF Routing

In this task, you will enable OSPF routing between PE and CE routers.

Activity Procedure

Complete these steps:

- Step 1** On the pod PE router, disable both interfaces toward the core devices (P1 and P2 routers).
- Step 2** On the pod CE router, enable OSPF area 0 on the Loopback0 and first Gigabit Ethernet interfaces. Use OSPF process ID 1.
- Step 3** On the pod PE router, enable OSPF area 0 on the Loopback0 interface and the first and second Gigabit Ethernet interfaces. Use OSPF process ID 1.
- Step 4** On the pod CE router, disable BGP neighbors (IPv4 and IPv6).
- Step 5** On the pod CE router, verify that PIM-SM is enabled on the Loopback0 interface and the first Gigabit Ethernet interface. On the pod PE router, verify that PIM-SM is enabled on the Loopback0 interface and on the first and second Gigabit Ethernet interfaces.

The CE router (Cisco IOS) output should be similar to the following, taken from Pod 3:

```
CE3# show running-config interface Loopback 0
Building configuration...

Current configuration : 153 bytes
!
interface Loopback0
 ip address 10.3.10.1 255.255.255.255
 ip router isis
 ip pim sparse-mode
 ipv6 address 2001:DB8:10:3:10::1/128
 ipv6 enable
end

CE3# show running-config interface GigabitEthernet 0/0
Building configuration...

Current configuration : 193 bytes
!
interface GigabitEthernet0/0
 ip address 192.168.103.31 255.255.255.0
 ip router isis
 ip pim sparse-mode
 duplex full
 speed 100
 ipv6 address 2001:DB8:192:168:103::31/80
 ipv6 enable
end
```

The PE router (Cisco IOS XR) output should be similar to the following, taken from Pod 3:

```
RP/0/RSP0/CPU0:PE3# show running-config router pim
router pim
 address-family ipv4
  rp-address 10.0.1.1
  interface Loopback0
  enable
  !
  interface GigabitEthernet0/0/0/0
  enable
  !
  interface GigabitEthernet0/0/0/1
  enable
  !
  interface GigabitEthernet0/0/0/2
  enable
  !
  interface GigabitEthernet0/0/0/3
  enable
  !
  !
  !
```

The CE router (Cisco IOS) output should be similar to the following, taken from Pod 4:

```
CE4# show running-config interface Loopback 0
Building configuration...

Current configuration : 153 bytes
!
interface Loopback0
 ip address 10.4.10.1 255.255.255.255
 ip router isis
 ip pim sparse-mode
 ipv6 address 2001:DB8:10:4:10::1/128
 ipv6 enable
end

CE4# show running-config interface GigabitEthernet 0/0
Building configuration...

Current configuration : 193 bytes
!
interface GigabitEthernet0/0
 ip address 192.168.104.41 255.255.255.0
 ip router isis
 ip pim sparse-mode
 duplex full
 speed 100
 ipv6 address 2001:DB8:192:168:104::41/80
 ipv6 enable
end
```

The PE router (Cisco IOS XE) output should be similar to the following, taken from Pod 4:

```

PE4# show running-config interface Loopback 0
Building configuration...
Current configuration : 170 bytes
!
interface Loopback0
 ip address 10.4.1.1 255.255.255.255
 ip router isis
 ip pim sparse-mode
 ipv6 address 2001:DB8:10:4:1::1/128
 ipv6 enable
 ipv6 router isis
end

PE4# show running-config interface GigabitEthernet 0/0/0
Building configuration...

Current configuration : 359 bytes
!
interface GigabitEthernet0/0/0
 ip address 192.168.104.40 255.255.255.0
 ip router isis
 ip pim sparse-mode
 ip verify unicast source reachable-via rx
 ip igmp query-interval 30
 speed 100
 no negotiation auto
 ipv6 address 2001:DB8:192:168:104::40/80
 ipv6 enable
 ipv6 mld query-interval 60
 ipv6 router isis
 cdp enable
 isis circuit-type level-1
end

PE4# show running-config interface GigabitEthernet 0/0/1
Building configuration...

Current configuration : 240 bytes
!
interface GigabitEthernet0/0/1
 ip address 192.168.134.40 255.255.255.0
 ip router isis
 ip pim sparse-mode
 negotiation auto
 ipv6 address 2001:DB8:192:168:134::40/80
 ipv6 enable
 ipv6 router isis
 isis circuit-type level-2-only
end

```

Step 6 On the pod CE and PE routers, remove static RP configuration.

Step 7 Simulate multicast receivers for group 224.1.x.1 (or 224.1.y.1) on the Loopback0 interface on the CE and PE routers. The routers should be interested only in traffic coming from the other pod CE router GigabitEthernet0/0 interface.

NOTE: Throughout the lab exercise, use the **ip igmp join-group** command on Cisco IOS and IOS-XE Software under interface configuration mode to simulate multicast receivers. On the Cisco IOS XR Software, use the **join-group** command under interface configuration mode under router IGMP configuration mode.

Verification

You have completed this task when you attain these results:

Step 1 On the pod CE router, verify the IP routing table. You should see the Loopback 0 subnets of the other pod PE and CE router.

The CE router (Cisco IOS) output should be similar to the following, taken from Pod 3:

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

Gateway of last resort is not set

10.0.0.0/32 is subnetted, 5 subnets
O    10.3.1.1 [110/2] via 192.168.103.30, 00:29:08, GigabitEthernet0/0
C    10.3.10.1 is directly connected, Loopback0
C    10.3.100.1 is directly connected, Loopback1
O    10.4.1.1 [110/3] via 192.168.103.30, 00:29:08, GigabitEthernet0/0
O    10.4.10.1 [110/4] via 192.168.103.30, 00:00:52, GigabitEthernet0/0
192.168.103.0/24 is variably subnetted, 2 subnets, 2 masks
C    192.168.103.0/24 is directly connected, GigabitEthernet0/0
L    192.168.103.31/32 is directly connected, GigabitEthernet0/0
O    192.168.104.0/24
      [110/3] via 192.168.103.30, 00:29:08, GigabitEthernet0/0
O    192.168.134.0/24
      [110/2] via 192.168.103.30, 00:29:08, GigabitEthernet0/0
```

The CE router (Cisco IOS) output should be similar to the following, taken from Pod 4:

```
CE4# show ip route
```

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

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

```
10.0.0.0/32 is subnetted, 5 subnets  
O    10.3.1.1 [110/3] via 192.168.104.40, 00:32:16, GigabitEthernet0/0  
O    10.3.10.1 [110/4] via 192.168.104.40, 00:30:27, GigabitEthernet0/0  
O    10.4.1.1 [110/2] via 192.168.104.40, 00:32:48, GigabitEthernet0/0  
C    10.4.10.1 is directly connected, Loopback0  
C    10.4.100.1 is directly connected, Loopback1  
O    192.168.103.0/24  
      [110/3] via 192.168.104.40, 00:32:16, GigabitEthernet0/0  
192.168.104.0/24 is variably subnetted, 2 subnets, 2 masks  
C    192.168.104.0/24 is directly connected, GigabitEthernet0/0  
L    192.168.104.41/32 is directly connected, GigabitEthernet0/0  
O    192.168.134.0/24  
      [110/2] via 192.168.104.40, 00:32:16, GigabitEthernet0/0
```

Task 2: Enable BSR

In this task, you will configure and verify BSR.

Activity Procedure

Complete this step:

Step 1 On the pod PE router, use the Loopback0 interface to configure a BSR candidate and BSR RP candidate.

Verification

You have completed this task when you attain these results:

- Step 1** Before you continue with the verification, clear the group-to-RP mapping table on the pod CE and PE router. The **clear ip pim rp-mapping** command should be entered on the Cisco IOS/IOS XE router and the **clear pim autorp** command should be entered on the Cisco IOS XR router.
- Step 2** On the pod PE router, verify that group-to-RP mapping information was obtained from the BSR. The PE router (Cisco IOS XR) output should be similar to the following, taken from Pod 3:

```
RP/0/RSP0/CPU0:PE3# show pim group-map
```

IP PIM Group Mapping Table

(* indicates group mappings being used)

(+ indicates BSR group mappings active in MRIB)

Group Range	Proto	Client	Groups	RP address	Info
224.0.1.39/32*	DM	perm	0	0.0.0.0	
224.0.1.40/32*	DM	perm	1	0.0.0.0	
224.0.0.0/24*	NO	perm	0	0.0.0.0	
232.0.0.0/8*	SSM	config	0	0.0.0.0	
224.0.0.0/4*	SM	bsr+	1	10.4.1.1	RPF: Gi0/0/0/1, 192.168.134.40
224.0.0.0/4	SM	bsr	0	10.3.1.1	RPF: Null, 0.0.0.0
224.0.0.0/4	SM	static	0	0.0.0.0	RPF: Null, 0.0.0.0

The PE router (Cisco IOS XE) output should be similar to the following, taken from Pod 4:

```
PE4# show ip pim rp mapping
```

PIM Group-to-RP Mappings

This system is a candidate RP (v2)

Group(s) 224.0.0.0/4

RP 10.4.1.1 (?), v2

Info source: 10.3.1.1 (?), via bootstrap, priority 0, holdtime 150

Uptime: 00:02:03, expires: 00:01:39

RP 10.3.1.1 (?), v2

Info source: 10.3.1.1 (?), via bootstrap, priority 192, holdtime 150

Uptime: 00:02:03, expires: 00:01:39

Step 3 On the pod CE router, verify that group-to-RP mapping information was obtained from the BSR. All routers in the team should have same group-to-RP mapping information.

In the output, the RP with IP address 10.y.1.1 and better priority 0 is mapped to the 224.0.0.0/4 multicast groups.

The CE router (Cisco IOS) output should be similar to the following, taken from Pod 3:

```
CE3# show ip pim rp mapping
```

PIM Group-to-RP Mappings

Group(s) 224.0.0.0/4

RP 10.4.1.1 (?), v2

Info source: 10.3.1.1 (?), via bootstrap, priority 0, holdtime 150

Uptime: 00:04:19, expires: 00:02:22

RP 10.3.1.1 (?), v2

Info source: 10.3.1.1 (?), via bootstrap, priority 192, holdtime 150

Uptime: 00:04:19, expires: 00:02:22

The CE router (Cisco IOS) output should be similar to the following, taken from Pod 4:

```
CE4# show ip pim rp mapping
PIM Group-to-RP Mappings
```

```
Group(s) 224.0.0.0/4
  RP 10.4.1.1 (?), v2
    Info source: 10.3.1.1 (?), via bootstrap, priority 0, holdtime 150
    Uptime: 00:04:16, expires: 00:02:26
  RP 10.3.1.1 (?), v2
    Info source: 10.3.1.1 (?), via bootstrap, priority 192, holdtime 150
    Uptime: 00:04:16, expires: 00:02:26
```

- Step 4** On the pod CE router, verify that the RP with the better (lower number) priority is used in multicast routing.

The CE router (Cisco IOS) output should be similar to the following, taken from Pod 3:

```
CE3# show ip mroute | include RP
  L - Local, P - Pruned, R - RP-bit set, F - Register flag,
(*, 224.1.3.1), 00:18:05/00:02:16, RP 10.4.1.1, flags: SJCL
  Incoming interface: GigabitEthernet0/0, RPF nbr 192.168.103.30
<... output omitted ...>
```

The CE router (Cisco IOS) output should be similar to the following, taken from Pod 4:

```
CE4# show ip mroute | include RP
  L - Local, P - Pruned, R - RP-bit set, F - Register flag,
(*, 224.1.4.1), 00:17:12/00:02:21, RP 10.4.1.1, flags: SJCL
  Incoming interface: GigabitEthernet0/0, RPF nbr 192.168.104.40
<... output omitted ...>
```

- Step 5** From the neighbor pod CE router, ping to the multicast group configured on your pod CE and PE routers. Use neighbor pod CE router Gigabit Ethernet 0/0 interface as a source.

Students in Pod 3 will start the ping from the CE4 router. The CE4 router (Cisco IOS) output should be similar to the following, taken from Pod 3.

```
CE4# ping 224.1.3.1 repeat 10 source GigabitEthernet 0/0
Type escape sequence to abort.
Sending 10, 100-byte ICMP Echos to 224.1.3.1, timeout is 2 seconds:
Packet sent with a source address of 192.168.104.41

Reply to request 0 from 10.3.10.1, 1 ms
Reply to request 0 from 10.3.1.1, 4 ms
Reply to request 1 from 10.3.10.1, 1 ms
Reply to request 1 from 10.3.1.1, 1 ms
Reply to request 2 from 10.3.10.1, 1 ms
Reply to request 2 from 10.3.1.1, 1 ms
```

Students in Pod 4 will start the ping from the CE3 router. The CE3 router (Cisco IOS) output should be similar to the following, taken from Pod 4.

```
CE3# ping 224.1.4.1 repeat 10 source GigabitEthernet 0/0
Type escape sequence to abort.
Sending 10, 100-byte ICMP Echos to 224.1.4.1, timeout is 2 seconds:
Packet sent with a source address of 192.168.103.31
```

```
Reply to request 0 from 10.4.1.1, 1 ms
Reply to request 0 from 10.4.10.1, 1 ms
Reply to request 1 from 10.4.1.1, 1 ms
Reply to request 1 from 10.4.10.1, 1 ms
Reply to request 2 from 10.4.10.1, 1 ms
Reply to request 2 from 10.4.1.1, 1 ms
```

Task 3: Enable Anycast RP


In this task, you will configure and verify Anycast RP.

Activity Procedure

Complete these steps:

- Step 1** On the pod PE router, remove the BSR candidate and BSR RP candidate configuration.
- Step 2** On the pod PE router, enable the Loopback1 interface and assign IP address 1.1.1.1/32. Make sure that the Loopback1 subnet is announced to the pod CE router via the OSPF routing protocol. On the newly configured Loopback1 interface, enable PIM-SM.
- Step 3** On the pod CE and PE routers, configure 1.1.1.1 as a static RP.
- Step 4** Establish MSDP adjacency between the pod PE and neighbor pod PE routers. Use the PE router Loopback0 IP address as the source address of MSDP packets and as the originator ID.

The figure shows what you will accomplish in this task.



Verification

You have completed this task when you attain these results:

- Step 1** Before you continue with the verification, clear the group-to-RP mapping table on the pod CE and PE routers. The **clear ip pim rp-mapping** command should be entered on the Cisco IOS/IOS XE router, and the **clear pim bsr** command should be entered on the Cisco IOS XR router.
- Step 2** On the pod PE router, verify that an MSDP session was established with the neighbor pod PE router.

The PE router (Cisco IOS XR) output should be similar to the following, taken from Pod 3:

```
RP/0/RSP0/CPU0:PE3# show msdp peer
MSDP Peer 10.4.1.1 (?), AS 0
Description:
Connection status:
  State: Up, Resets: 0, Connection Source: 10.3.1.1
  Uptime (Downtime): 00:02:35, SA messages received: 0
  TLV messages sent/received: 6/2
Output messages discarded: 0
  Connection and counters cleared 00:02:35 ago
SA Filtering:
  Input (S,G) filter: none
  Input RP filter: none
  Output (S,G) filter: none
  Output RP filter: none
SA-Requests:
  Input filter: none
  Sending SA-Requests to peer: disabled
Password: None
Peer ttl threshold: 2
Input queue size: 0, Output queue size: 0
KeepAlive timer period: 30
Peer Timeout timer period: 75
```

The PE router (Cisco IOS XE) output should be similar to the following, taken from Pod 4:

```

PE4# show ip msdp peer
MSDP Peer 10.3.1.1 (?), AS ?
  Connection status:
    State: Up, Resets: 0, Connection source: Loopback0 (10.4.1.1)
    Uptime(Downtime): 00:03:19, Messages sent/received: 3/7
    Output messages discarded: 0
    Connection and counters cleared 00:03:25 ago
  SA Filtering:
    Input (S,G) filter: none, route-map: none
    Input RP filter: none, route-map: none
    Output (S,G) filter: none, route-map: none
    Output RP filter: none, route-map: none
  SA-Requests:
    Input filter: none
  Peer ttl threshold: 0
  SAs learned from this peer: 0
  Number of connection transitions to Established state: 1
    Input queue size: 0, Output queue size: 0
  MD5 signature protection on MSDP TCP connection: not enabled
  Message counters:
    RPF Failure count: 0
    SA Messages in/out: 0/0
    SA Requests in: 0
    SA Responses out: 0
    Data Packets in/out: 0/0

```

Step 3 From the neighbor pod CE router, ping to the multicast group configured on your pod CE and PE routers. Use the neighbor pod CE router Gigabit Ethernet 0/0 interface as a source.

Students in Pod 3 will start the ping from the CE4 router. The CE4 router (Cisco IOS) output should be similar to the following, taken from Pod 3.

```

CE4# ping 224.1.3.1 repeat 10 source GigabitEthernet 0/0
Type escape sequence to abort.
Sending 10, 100-byte ICMP Echos to 224.1.3.1, timeout is 2 seconds:
Packet sent with a source address of 192.168.104.41

Reply to request 0 from 10.3.1.1, 1 ms
Reply to request 0 from 10.3.10.1, 1 ms
Reply to request 1 from 10.3.10.1, 1 ms
Reply to request 1 from 10.3.1.1, 1 ms
Reply to request 2 from 10.3.10.1, 1 ms
Reply to request 2 from 10.3.1.1, 1 ms
Reply to request 3 from 10.3.10.1, 1 ms
Reply to request 3 from 10.3.1.1, 1 ms

```

Students in Pod 4 will start the ping from the CE3 router. The CE3 router (Cisco IOS) output should be similar to the following, taken from Pod 4.

```

CE3# ping 224.1.4.1 repeat 10 source GigabitEthernet 0/0
Type escape sequence to abort.
Sending 10, 100-byte ICMP Echos to 224.1.4.1, timeout is 2 seconds:
Packet sent with a source address of 192.168.103.31

Reply to request 0 from 10.4.1.1, 1 ms
Reply to request 0 from 10.4.10.1, 1 ms
Reply to request 1 from 10.4.1.1, 1 ms
Reply to request 1 from 10.4.10.1, 1 ms
Reply to request 2 from 10.4.1.1, 1 ms
Reply to request 2 from 10.4.10.1, 1 ms

```

Step 4 Verify the MSDP SA cache on the PE routers:

The PE router (Cisco IOS XR) output should be similar to the following, taken from Pod 3:

```

RP/0/RSP0/CPU0:PE3# show msdp sa-cache

MSDP Flags:
E - set MRIB E flag , L - domain local source is active,
EA - externally active source, PI - PIM is interested in the group,
DE - SAs have been denied. Timers age/expiration,
Cache Entry:
(192.168.104.41, 224.1.3.1), RP 10.4.1.1, MBGP/AS 0, 00:00:09/00:02:20
  Learned from peer 10.4.1.1, RPF peer 10.4.1.1
  SAs recvd 1, Encapsulated data received: 100
  grp flags: PI, src flags: E, EA, PI
(192.168.103.31, 224.1.4.1), RP 1.1.1.1, MBGP/AS 0, 00:00:04/local
  Learned from peer local, RPF peer local
  SAs recvd 0, Encapsulated data received: 0
  grp flags: none, src flags: L

```

The PE router (Cisco IOS XE) output should be similar to the following, taken from Pod 4:

```

PE4# show ip msdp sa-cache
MSDP Source-Active Cache - 2 entries
(192.168.103.31, 224.1.3.1), RP 10.3.1.1, BGP/AS 0, 00:02:09/00:05:43, Peer 10.3.1.1
(192.168.103.31, 224.1.4.1), RP 10.3.1.1, BGP/AS 0, 00:01:38/00:05:43, Peer 10.3.1.1

```


Hardware Lab 8: Implement a DHCPv6 Server with Prefix Delegation

Activity Objective

In this lab activity, you will configure the PE router to act as a prefix delegation DHCPv6 server. You will configure the PE router to delegate a prefix to the CE router. You will then configure the CE router interface with an IPv6 address from the delegated prefix. You will also configure the CE router to act as a DHCPv6 Lite server, with DNS server IP address options obtained from the PE router. This scenario is valid, where the service provider would assign a prefix to a CE router using prefix delegation, and the CE router would then advertise the prefix to LAN endpoints to enable stateless autoconfiguration.

This lab exercise will be performed only in the pod with Cisco IOS XE router.


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

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

- Implement a prefix delegation DHCPv6 server
- Implement a DHCPv6 client
- Implement a DHCPv6 Lite server

Visual Objective

The figure illustrates what you will accomplish in this activity.



Command List

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

Cisco IOS/IOS XE Commands

Command	Description
configure terminal	Enters configuration mode.
dns-server <i>IPv6_address</i>	Specifies the DNS IPv6 servers available to a DHCP server for the IPv6 client.
import dns-server	Imports the DNS name server option into a DHCP server for the IPv6 client.
interface <i>interface</i>	Enters interface configuration mode.
ipv6 address autoconfig	Enables automatic configuration of IPv6 addresses using stateless autoconfiguration on an interface and enables IPv6 processing on the interface.
ipv6 address <i>prefix_name suffix</i> <i>prefix-length</i>	Configures an IPv6 address based on an IPv6 general prefix and enables IPv6 processing on an interface.

Command	Description
ipv6 dhcp client pd <i>prefix_name</i>	Enables a DHCP for IPv6 client process and enables requests for prefix delegation through a specified interface.
ipv6 dhcp pool <i>pool_name</i>	Configures a DHCP for IPv6 server configuration information pool and enters DHCP for IPv6 pool configuration mode.
ipv6 dhcp server <i>pool_name</i>	Enables a DHCP server for IPv6 on an interface.
ipv6 local pool <i>prefix_pool_name prefix_length assigned_length</i>	Configures a local IPv6 prefix pool.
ipv6 nd other-config-flag	Sets the "other stateful configuration" flag in IPv6 router advertisements.
ping <i>dest_ip_source source_interface</i>	Verifies connectivity between the source IP and the destination IP.
prefix-delegation pool <i>prefix_pool_name</i>	Specifies a named IPv6 local prefix pool from which prefixes are delegated to the DHCP server for IPv6 clients.
show ipv6 dhcp interface	Displays DHCP for IPv6 interface information.
show ipv6 dhcp pool	Displays DHCP for IPv6 configuration pool information.
show ipv6 interface brief	Displays IPv6 addresses on interfaces and the status of interfaces.

Task 1: Configure a Prefix Delegation DHCPv6 Server and Client

In this task, you will configure a prefix delegation DHCPv6 server on the PE router (IOS XE only).

Activity Procedure

Complete these steps:

- Step 1** From the PE router (IOS XE only), ping the directly connected interface of the CE router using IPv6. The ping should be successful.

The PE router (Cisco IOS XE) output should be similar to the following, taken from Pod 4:

```
PE4# ping 2001:DB8:192:168:104::41
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:192:168:104::41, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

- Step 2** Configure the PE router (IOS XE only) as a prefix delegation DHCPv6 server with the following parameters:

Pod	Delegated Prefix	DNS Server
Pod y	2001:db8:100:y::/64	2001:db8:100::y

y should be substituted with 2 or 4, depending on your team assignment.

Enable the DHCPv6 server on the interface that is facing the CE router.

- Step 3** On the CE router, enable a DHCPv6 prefix delegation client on the interface that is facing the PE router (IOS XE only). Use SP_ASSIGNED_PREFIX as the name of the delegated prefix.

- Step 4** On the CE router assign the first IP address from the delegated prefix to the GigabitEthernet0/1 interface.

Verification

You have completed this task when you attain these results:

- Step 1** Verify the assigned IPv6 address on the CE GigabitEthernet0/1 interface.

The CE router (Cisco IOS) output should be similar to the following, taken from Pod 4:

```
CE4# show ipv6 interface brief
En0/0                                [administratively down/down]
    unassigned
GigabitEthernet0/0                  [up/up]
    FE80::5257:A8FF:FE4E:D548
    2001:DB8:192:168:104::41
GigabitEthernet0/1                  [up/up]
    FE80::5257:A8FF:FE4E:D549
    2001:DB8:100:4::1
Loopback0                           [up/up]
    FE80::5257:A8FF:FE4E:D548
    2001:DB8:10:4:10::1
Loopback1                           [up/up]
    unassigned
Tunnel0                             [up/up]
    unassigned
```

- Step 2** From the CE router, ping a nonexistent hostname.

The CE router (Cisco IOS) output should be similar to the following, taken from Pod 4:

```
CE4# ping asfdg
Translating "asfdg"...domain server (2001:DB8:100::4)

Translating "asfdg"...domain server (2001:DB8:100::4)
<... output omitted ...>
```

- Step 3** On the PE router, verify DHCP bindings.

The PE router (Cisco IOS XE) output should be similar to the following, taken from Pod 4:

```

PE4# show ipv6 dhcp binding
Client: FE80::5257:A8FF:FE4E:D548
DUID: 000300015057A84ED548
Username : unassigned
VRF : default
Interface : GigabitEthernet0/0/0
IA PD: IA ID 0x00040001, T1 302400, T2 483840
Prefix: 2001:DB8:100:4::/64
        preferred lifetime 604800, valid lifetime 2592000
        expires at Aug 01 2014 07:24 PM (2586754 seconds)

```

Task 2: Configure DHCPv6 Lite Server

In this task, you will configure the CE router, which also acts as a DHCPv6 client, as a DHCPv6 Lite server. The CE router will advertise the DNS server IP address, which will be obtained from the PE router, to LAN clients. LAN clients will obtain IPv6 addresses using stateless autoconfiguration.

Activity Procedure

Complete these steps:

- Step 1** On the CE router, configure a DHCPv6 pool. The DNS server should be imported as received from the PE router (IOS XE only).
- Step 2** On the CE router, enable the DHCPv6 server on the GigabitEthernet0/1 interface. Configure the router to instruct DHCP clients not to use DHCP for address assignments. However, clients should obtain other parameters, such as the DNS server, using DHCP.
- Step 3** From the CE router, remove the previously configured IPv6 address on the GigabitEthernet0/1 interface.
- Step 4** Optionally, on the CE router, configure the GigabitEthernet0/1 interface to obtain an IP address via stateless autoconfiguration.

Verification

You have completed this task when you attain these results:

- Step 1** On the CE router, verify the configured DHCP pools.

The CE router (Cisco IOS) output should be similar to the following, taken from Pod 4:

```

CE4# show ipv6 dhcp pool
DHCPv6 pool: POOL
Imported DNS server: 2001:DB8:100::4
Active clients: 0

```

- Step 2** On the CE router, verify how DHCP is enabled on interfaces. You should see that one interface acts as the DHCP client and that the other acts as the DHCP server.

The CE router (Cisco IOS) output should be similar to the following, taken from Pod 4:

```

CE4# show ipv6 dhcp interface
GigabitEthernet0/0 is in client mode
Prefix State is OPEN
Renew will be sent in 3d11h
Address State is IDLE
List of known servers:
  Reachable via address: FE80::462B:3FF:FE53:FC00
  DUID: 00030001442B0353FC00
  Preference: 0
Configuration parameters:
  IA PD: IA ID 0x00040001, T1 302400, T2 483840
  Prefix: 2001:DB8:100:4::/64
           preferred lifetime 604800, valid lifetime 2592000
           expires at Aug 02 2014 08:29 AM (2591835 seconds)
  DNS server: 2001:DB8:100::4
  Information refresh time: 0
Prefix name: SP ASSIGNED PREFIX
Prefix Rapid-Commit: disabled
Address Rapid-Commit: disabled
GigabitEthernet0/1 is in server mode
Using pool: POOL
Preference value: 0
Hint from client: ignored
Rapid-Commit: disabled

```

Step 3 From the CE router, ping a nonexistent hostname. This action will trigger DNS name lookup and you will be able to verify whether a DNS server has been assigned via DHCP.

The CE router (Cisco IOS) output should be similar to the following, taken from Pod 4:

```

CE4# ping wf
Translating "wf"...domain server (2001:DB8:100::4)

Translating "wf"...domain server (2001:DB8:100::4)
<... output omitted ...>

```

Hardware Lab 9: Implement IPv6 Multicasting

Activity Objective

In this lab activity, you will first implement IPv6 multicast using embedded RPs. Your pod CE router will act as the multicast receiver, and the other pod CE router will act as the multicast source. Your pod PE router will act as the RP.


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

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 this objective:

- Implement IPv6 multicast using embedded RPs

Visual Objective

The figure illustrates what you will accomplish in this activity.



Command List

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

Cisco IOS/IOS XE Commands

Command	Description
configure terminal	Enters configuration mode.
interface interface	Enters interface configuration mode.
ipv6 address ipv6_address/prefix	Assigns an IPv6 address to an interface.
ipv6 igmp join-group group_address	Configures an interface on the router to join the specified group or channel.
ipv6 multicast-routing	Enables IPv6 multicast routing.
ipv6 pim rp-address RP_address	Statically configures the address of a PIM RP for multicast groups.
ipv6 pim spt-threshold threshold	Configures when a PIM leaf router should join the shortest path source tree.

Command	Description
ipv6 ospf <i>process-name area area-id</i>	Enables the OSPFv3 routing protocol on an interface and advertises a network on the interface.
ping <i>dest_IP source source_IP</i>	Verifies connectivity between the source IP and the destination IP (IPv4 and IPv6).
show ip protocols	Displays configured routing protocols for IPv4.
show ipv6 mroute	Displays the contents of the multicast routing (mroute) table.
show ipv6 pim interface	Displays information about interfaces that are configured for PIM.
show ipv6 pim neighbor	Displays information about PIM neighbors discovered by PIMv1 router query messages or PIMv2 hello messages.
show ipv6 protocols	Displays configured routing protocols for IPv6.
show ipv6 route	Displays a routing table on a router.
show ipv6 ospf neighbor interface	Displays OSPFv3 neighbors or OSPF enabled interfaces.

Cisco IOS XR Commands

Command	Description
address-family ipv6	Enters an IPv6 address family under a specific configuration mode. Enables an interface for a specified address family under IS-IS configuration mode
commit	Commits changes to the running configuration.
configure	Enters configuration mode.
embedded-rp <i>RP_address access_list_name</i>	Statically configures the address of an embedded RP for multicast groups under PIM configuration mode.
enable	Enables an interface for multicast routing or PIM (under the appropriate configuration mode).
interface <i>interface</i>	Enters interface configuration mode.
ipv6 access-list <i>acl_name</i>	Creates an ACL and enters access list configuration mode.
multicast-routing	Enters multicast routing configuration mode.

Command	Description
<code>[permit deny] protocol source_address destination_address</code>	Creates an entry in an ACL.
<code>ping dest_IP source source_IP</code>	Verifies connectivity between the source IP and the destination IP (IPv4 and IPv6).
<code>router ospfv3 process-name</code>	Enters OSPFv3 configuration mode.
<code>router pim</code>	Enters PIM configuration mode.
<code>show mrib ipv6 route</code>	Displays the contents of the multicast routing (mroute) table.
<code>show pim ipv6 interface</code>	Displays information about interfaces that are configured for PIM.
<code>show pim ipv6 neighbor</code>	Displays information about PIM neighbors discovered by PIMv1 router query messages or PIMv2 hello messages.
<code>spt-threshold infinity</code>	Configures a PIM leaf router to join the shortest path source tree immediately.

Task 1: Create a New Loopback Interface and Verify Connectivity

In this task, you will configure a new loopback interface on the PE router that will act as the RP for your pod.

Activity Procedure

Complete these steps:

- Step 1** On the PE router, create a new loopback interface. Use x0 (or y0) as the interface identifier. This interface will be used as the RP for your pod.
- Step 2** On the PE router, assign the following IPv6 address to the loopback interface x0 or the loopback y0 interface:

Router	IPv6 Address
PE _x	2001.db8:x:x::1/128
PE _y	2001.db8:y:y::1/128

- Step 3** On the PE router advertise the loopback x0 interface or the loopbacky0 interface using the OSPFv3 routing protocol into area 0.
- Step 4** On the CE router, verify that the GigabitEthernet0/0 and Loopback 0 interfaces are enabled for OSPFv3 in the area 0.

The CE router (Cisco IOS) output should be similar to the following, taken from Pod 3:

```
CE3# show ipv6 ospf interface
Loopback0 is up, line protocol is up
  Link Local Address FE80::A64C:11FF:FE04:4C98, Interface ID 10
  Area 0, Process ID 1, Instance ID 0, Router ID 10.3.100.1
  Network Type LOOPBACK, Cost: 1
  Loopback interface is treated as a stub Host
GigabitEthernet0/0 is up, line protocol is up
  Link Local Address FE80::A64C:11FF:FE04:4C98, Interface ID 4
  Area 0, Process ID 1, Instance ID 0, Router ID 10.3.100.1
  Network Type BROADCAST, Cost: 1
  Transmit Delay is 1 sec, State DR, Priority 1
  Designated Router (ID) 10.3.100.1, local address FE80::A64C:11FF:FE04:4C98
  Backup Designated router (ID) 10.3.1.1, local address FE80::6E9C:EDFF:FE22:55C8
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:05
  Graceful restart helper support enabled
  Index 1/1/1, flood queue length 0
  Next 0x0(0)/0x0(0)/0x0(0)
  Last flood scan length is 1, maximum is 3
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 10.3.1.1 (Backup Designated Router)
  Suppress hello for 0 neighbor(s)
```

The CE router (Cisco IOS) output should be similar to the following, taken from Pod 4:

```
CE4# show ipv6 ospf interface
Loopback0 is up, line protocol is up
  Link Local Address FE80::5257:A8FF:FE4E:D548, Interface ID 10
  Area 0, Process ID 1, Instance ID 0, Router ID 10.4.100.1
  Network Type LOOPBACK, Cost: 1
  Loopback interface is treated as a stub Host
GigabitEthernet0/0 is up, line protocol is up
  Link Local Address FE80::5257:A8FF:FE4E:D548, Interface ID 4
  Area 0, Process ID 1, Instance ID 0, Router ID 10.4.100.1
  Network Type BROADCAST, Cost: 1
  Transmit Delay is 1 sec, State DR, Priority 1
  Designated Router (ID) 10.4.100.1, local address FE80::5257:A8FF:FE4E:D548
  Backup Designated router (ID) 10.4.1.1, local address FE80::462B:3FF:FE53:FC00
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:03
  Graceful restart helper support enabled
  Index 1/1/1, flood queue length 0
  Next 0x0(0)/0x0(0)/0x0(0)
  Last flood scan length is 0, maximum is 4
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 10.4.1.1 (Backup Designated Router)
  Suppress hello for 0 neighbor(s)
```

- Step 5** On the CE router, verify that the OSPFv3 adjacency is established with the PE router.

The CE router (Cisco IOS) output should be similar to the following, taken from Pod 3:

```
CE3# show ipv6 ospf neighbor
```

```
OSPFv3 Router with ID (10.3.100.1) (Process ID 1)
```

Neighbor ID	Pri	State	Dead Time	Interface ID	Interface
10.3.1.1	1	FULL/BDR	00:00:31	12	GigabitEthernet0/0

The CE router (Cisco IOS) output should be similar to the following, taken from Pod 4:

```
CE4# show ipv6 ospf neighbor
```

```
OSPFv3 Router with ID (10.4.100.1) (Process ID 1)
```

Neighbor ID	Pri	State	Dead Time	Interface ID	Interface
10.4.1.1	1	FULL/BDR	00:00:30	8	GigabitEthernet0/0

Step 6 From the CE router Loopback 0, verify the IPv6 connectivity to the PE router Loopback x0 interface or the Loopback y0 interface. The ping should be successful.

The CE router (Cisco IOS) output should be similar to the following, taken from Pod 3:

```
CE3# ping 2001:db8:3:3::1 source Loopback 0
```

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

```
Sending 5, 100-byte ICMP Echos to 2001:DB8:3:3::1, timeout is 2 seconds:
```

```
Packet sent with a source address of 2001:DB8:10:3:10::1
```

```
!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/8 ms
```

The CE router (Cisco IOS) output should be similar to the following, taken from Pod 4:

```
CE4# ping 2001:DB8:4:4::1 source Loopback 0
```

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

```
Sending 5, 100-byte ICMP Echos to 2001:DB8:4:4::1, timeout is 2 seconds:
```

```
Packet sent with a source address of 2001:DB8:10:4:10::1
```

```
!!!!
```

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

Step 7 On the CE router, verify that the IPv6 address from the new Loopback x0 interface or the Loopback y0 interface on the PE has been learned through OSPFv3.

The CE router (Cisco IOS) output should be similar to the following, taken from Pod 3:

```
CE3# show ipv6 route 2001:db8:3:3::1
```

```
Routing entry for 2001:DB8:3:3::1/128
```

```
Known via "ospf 1", distance 110, metric 1, type intra area
```

```
Route count is 1/1, share count 0
```

```
Routing paths:
```

```
FE80::6E9C:EDFF:FE22:55C8, GigabitEthernet0/0
```

```
Last updated 00:28:42 ago
```

The CE router (Cisco IOS) output should be similar to the following, taken from Pod 4:

```
CE4# show ipv6 route 2001:DB8:4:4::1
Routing entry for 2001:DB8:4:4::1/128
  Known via "ospf 1", distance 110, metric 1, type intra area
  Route count is 1/1, share count 0
  Routing paths:
    FE80::462B:3FF:FE53:FC00, GigabitEthernet0/0
    Last updated 00:29:29 ago
```

Verification

No verification is needed for this task.

Task 2: Implement IPv6 Multicast Using Embedded RPs

In this task, you will implement IPv6 multicast on the PE router and CE routers. The pod PE router will act as the RP. The multicast group address will carry the RP IP address information.

Activity Procedure

Complete these steps:

- Step 1** On the CE router, enable IPv6 multicast routing. Configure the router to always stay on the shared tree.
- Step 2** On the PE router, enable IPv6 multicast routing. On the PE router running Cisco IOS XR Software, make sure that you enable PIM on all relevant interfaces, including the newly created loopback interface. Configure the router to always stay on the shared tree.
- Step 3** On the PE router, verify the PIM state for IPv6 on interfaces.
The PE router (Cisco IOS XR) output should be similar to the following, taken from Pod 3:

```
RP/0/RSP0/CPU0:PE3# show pim ipv6 interface
```

```
PIM interfaces in VRF default
```

Interface	PIM	Nbr Count	Hello Intvl	DR Prior
Loopback0	on	1	30	1
Primary Address : fe80::b42d:e8ff:fe56:d20e				
DR : this system				
Loopback30	on	1	30	1
Primary Address : fe80::b42d:e8ff:fe56:d20e				
DR : this system				
GigabitEthernet0/0/0/0	on	2	30	1
Primary Address : fe80::6e9c:edff:fe22:55c8				
DR : fe80::a64c:11ff:fe04:4c98				
GigabitEthernet0/0/0/1	on	2	30	1
Primary Address : fe80::6e9c:edff:fe22:55c9				
DR : this system				
GigabitEthernet0/0/0/2	off	0	30	1
Primary Address :				
DR : not elected				
GigabitEthernet0/0/0/3	off	0	30	1
Primary Address :				
DR : not elected				

The PE router (Cisco IOS XE) output should be similar to the following, taken from Pod 4:

```

PE4# show ipv6 pim interface
Interface          PIM    Nbr    Hello  DR
                  Count Intvl  Prior

Null0              off    0      30     1
  Address: FE80::1
  DR       : not elected
Loopback1          off    0      30     1
  Address: ::
  DR       : not elected
Gi0/0/0            on     1      30     1
  Address: FE80::462B:3FF:FE53:FC00
  DR       : FE80::5257:A8FF:FE4E:D548
Gi0/0/1            on     1      30     1
  Address: FE80::462B:3FF:FE53:FC01
  DR       : FE80::6E9C:EDFF:FE22:55C9
Gi0/0/2            off    0      30     1
  Address: ::
  DR       : not elected
Gi0/0/3            off    0      30     1
  Address: ::
  DR       : not elected
POS0/2/0           off    0      30     1
  Address: ::
  DR       : not elected
POS0/2/1           off    0      30     1
  Address: ::
  DR       : not elected
Loopback40         on     0      30     1
  Address: FE80::462B:3FF:FE53:FC00
  DR       : this system
Loopback0          on     0      30     1
  Address: FE80::462B:3FF:FE53:FC00
  DR       : this system
Tunnel0            off    0      30     1
  Address: FE80::462B:3FF:FE53:FC00
  DR       : not elected
Tunnel1            off    0      30     1
  Address: ::
  DR       : not elected
Tunnel2            off    0      30     1
  Address: ::
  DR       : not elected

```

Step 4 On the PE router, verify the PIM neighbors for IPv6.

The PE router (Cisco IOS XR) output should be similar to the following, taken from Pod 3:

```
RP/0/RSP0/CPU0:PE3# show pim ipv6 neighbor
```

PIM neighbors in VRF default

Flag: B - Bidir capable, P - Proxy capable, DR - Designated Router,
* indicates the neighbor created for this router

Loopback0

Neighbor Address	Uptime	Expires	DR	pri	DR	Flags
fe80::b42d:e8ff:fe56:d20e*	00:08:51	00:01:29	1		(DR)	B P

Loopback30

Neighbor Address	Uptime	Expires	DR	pri	DR	Flags
fe80::b42d:e8ff:fe56:d20e*	00:08:51	00:01:28	1		(DR)	B P

GigabitEthernet0/0/0/0

Neighbor Address	Uptime	Expires	DR	pri	DR	Flags
fe80::6e9c:edff:fe22:55c8*	1w2d	00:01:37	1			B P
fe80::a64c:11ff:fe04:4c98	00:10:44	00:01:29	1		(DR)	B

GigabitEthernet0/0/0/1

Neighbor Address	Uptime	Expires	DR	pri	DR	Flags
fe80::462b:3ff:fe53:fc01	00:08:45	00:01:24	1			B
fe80::6e9c:edff:fe22:55c9*	00:08:51	00:01:25	1		(DR)	B P

The PE router (Cisco IOS XE) output should be similar to the following, taken from Pod 4:

```
PE4# show ipv6 pim neighbor
```

PIM Neighbor Table

Mode: B - Bidir Capable, G - GenID Capable

Neighbor Address	Interface	Uptime	Expires	Mode	DR	pri
FE80::5257:A8FF:FE4E:D548	Gi0/0/0	00:12:11	00:01:28	B G	DR	1
FE80::6E9C:EDFF:FE22:55C9	Gi0/0/1	00:10:05	00:01:36	B G	DR	1

Step 5 Calculate the multicast group addresses that can be served by the RP assigned to your pod. Use site-local multicast group addresses. Complete the following table:

Router	RP Address	Group Addresses
PEx	2001.db8:x:x::1/128	
PEy	2001.db8:y:y::1/128	

Step 6 On the PE router, manually configure the RP address. Loopbackx0 (or Loopbacky0) will act as an RP for your pod.

On the PE router running Cisco IOS XR Software, you also must provide multicast groups that will be served by that RP using an ACL.

NOTE: Manual RP configuration on other routers is not needed because the other routers will learn the RP information from the multicast group IP address.

NOTE: Both pods from the same team should be finished with the previous steps at this point.

Step 7 How many multicast groups can an embedded RP serve?

Step 8 On the CE router enable Loopback 0 as a multicast receiver for one of the multicast groups that is served by your RP (for example, FF75:0140:2001:db8:x:x::1 or FF75:0140:2001:db8:y:y::1).

Note: Throughout the lab exercise, use the **ipv6 mld join-group** command on Cisco IOS and IOS-XE Software under interface configuration mode to simulate multicast receivers.

Step 9 Use Telnet to connect to the other pod CE router. Start the multicast traffic by pinging your multicast group address. Use the Loopback 0 interface as a source interface. Send a large number of ICMP packets—100, for example. You should see that your CE router replies to the ping.

Students in Pod 3 will start the ping from the CE4 router. The CE4 router (Cisco IOS) output should be similar to the following, taken from Pod 3.

```
CE4# ping FF75:140:2001:DB8:3:3:0:1 repeat 100 source 2001:DB8:10:4:10::1
Output Interface: Loopback0
Type escape sequence to abort.
Sending 100, 100-byte ICMP Echos to FF75:140:2001:DB8:3:3:0:1, timeout is 2 seconds:
Packet sent with a source address of 2001:DB8:10:4:10::1

Request 0 timed out
Request 1 timed out
Reply to request 2 received from 2001:DB8:10:3:10::1, 1 ms
Reply to request 3 received from 2001:DB8:10:3:10::1, 1 ms
Reply to request 4 received from 2001:DB8:10:3:10::1, 1 ms
Reply to request 5 received from 2001:DB8:10:3:10::1, 1 ms
Reply to request 6 received from 2001:DB8:10:3:10::1, 1 ms
<... output omitted ...>
```

Students in Pod 4 will start the ping from the CE3 router. The CE3 router (Cisco IOS) output should be similar to the following, taken from Pod 4.

```
CE3# ping FF75:140:2001:DB8:4:4:0:1 repeat 100 source 2001:DB8:10:3:10::1
Output Interface: Loopback0
Type escape sequence to abort.
Sending 100, 100-byte ICMP Echos to FF75:140:2001:DB8:4:4:0:1, timeout is 2 seconds:
Packet sent with a source address of 2001:DB8:10:3:10::1

Reply to request 0 received from 2001:DB8:10:4:10::1, 12 ms
Reply to request 1 received from 2001:DB8:10:4:10::1, 1 ms
Reply to request 2 received from 2001:DB8:10:4:10::1, 1 ms
<... output omitted ...>
```

Step 10 On the PE router examine the IPv6 multicast routing table.

The PE router (Cisco IOS XR) output should be similar to the following, taken from Pod 3:

```

RP/0/RSP0/CPU0:PE3# show mrib ipv6 route
<... output omitted ...>
(*,ff75:140:2001:db8:3:3::/96)
  RPF nbr: 2001:db8:3:3::1 Flags: L C P
  Up: 00:26:53
  Outgoing Interface List
    Decaps6tunnel0 Flags: NS DI, Up: 00:26:53

(*,ff75:140:2001:db8:3:3:0:1)
  RPF nbr: 2001:db8:3:3::1 Flags: C
  Up: 00:15:23
  Incoming Interface List
    Decaps6tunnel0 Flags: A, Up: 00:15:23
  Outgoing Interface List
    GigabitEthernet0/0/0/0 Flags: F NS, Up: 00:15:23

(2001:db8:10:4:10::1,ff75:140:2001:db8:3:3:0:1)
  RPF nbr: fe80::462b:3ff:fe53:fc01 Flags: L
  Up: 00:03:41
  Incoming Interface List
    GigabitEthernet0/0/0/1 Flags: A, Up: 00:03:41
  Outgoing Interface List
    GigabitEthernet0/0/0/0 Flags: F NS, Up: 00:03:41

(*,ff75:140:2001:db8:4:4::/96)
  RPF nbr: fe80::462b:3ff:fe53:fc01 Flags: C P
  Up: 00:01:51

(2001:db8:10:3:10::1,ff75:140:2001:db8:4:4:0:1)
  RPF nbr: fe80::a64c:11ff:fe04:4c98 Flags:
  Up: 00:01:51
  Incoming Interface List
    GigabitEthernet0/0/0/0 Flags: A, Up: 00:01:51
  Outgoing Interface List
    GigabitEthernet0/0/0/1 Flags: F NS, Up: 00:01:51
<... output omitted ...>

```

The PE router (Cisco IOS XE) output should be similar to the following, taken from Pod 4:

```

PE4# show ipv6 mroute
<... output omitted ...>
(2001:DB8:10:4:10::1, FF75:140:2001:DB8:3:3:0:1), 00:06:15/00:02:32, flags: ST
  Incoming interface: GigabitEthernet0/0/0
  RPF nbr: FE80::5257:A8FF:FE4E:D548
  Immediate Outgoing interface list:
    GigabitEthernet0/0/1, Forward, 00:06:15/00:02:32

(*, FF05::1:3), 1d02h/never, RP 2001:DB8:4:4::1, flags: SPCL
  Incoming interface: Tunnel4
  RPF nbr: 2001:DB8:4:4::1
  Immediate Outgoing interface list:
    GigabitEthernet0/0/0, Null, 1d02h/never

(*, FF75:140:2001:DB8:4:4:0:1), 00:10:37/00:02:52, RP 2001:DB8:4:4::1, flags: S
  Incoming interface: Tunnel4
  RPF nbr: 2001:DB8:4:4::1
  Immediate Outgoing interface list:
    GigabitEthernet0/0/0, Forward, 00:10:37/00:02:52

(2001:DB8:10:3:10::1, FF75:140:2001:DB8:4:4:0:1), 00:04:25/00:02:48, flags: ST
  Incoming interface: GigabitEthernet0/0/1
  RPF nbr: FE80::6E9C:EDFF:FE22:55C9
  Inherited Outgoing interface list:
    GigabitEthernet0/0/0, Forward, 00:10:37/00:02:52

```

Verification

No verification is needed for this task.

Hardware Lab 10: Implement Tunnels for IPv6

Activity Objective

In this lab activity, you will first establish a static IPv6-in-IPv4 tunnel between two CE routers in different pods in the same team. Then you will deploy 6RD tunnels and configure a 6RD CE router and 6RD Border Relay.


Note	Students from two different pods work as a team. The CE routers in both pods are running Cisco IOS Software. In this activity, students will work on CE routers. Students from both pods will work on the PE router from the second pod only. Students in the same team should coordinate their activities.
-------------	---

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

- Configure a static IPv6-in-IPv4 tunnel
- Configure 6RD

Visual Objective

The figure illustrates what you will accomplish in this activity.



Command List

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

Cisco IOS/IOS XE Commands

Command	Description
configure terminal	Enters configuration mode.
ipv6 address <i>IPv6_address</i>	Configures an IPv6 address on an interface.
interface <i>interface</i>	Enters interface configuration mode.
ipv6 enable	Enables an interface for IPv6.
interface tunnel <i>id</i>	Creates a tunnel interface and enters interface configuration mode.
tunnel source <i>interface</i>	Specifies a tunnel source.
tunnel destination <i>ip_address</i>	Specifies a tunnel destination.

Command	Description
tunnel mode ipv6ip	Sets tunnel mode to IPv6-in-IPv4.
tunnel mode ipv6ip 6rd	Sets tunnel mode to 6RD.
tunnel 6rd ipv4 prefix-len <i>length</i>	Sets a common prefix length for 6RD.
tunnel 6rd prefix <i>prefix</i>	Sets a 6RD prefix.
tunnel 6rd br	Sets a 6RD Border Relay IP address.
ipv6 route <i>prefix/length outgoing_interface next_hop_IP_address</i>	Configures a static IPv6 route.
ping <i>dest_ip source source_interface</i>	Verifies connectivity between the source IP and the destination IP.
router bgp <i>as_number</i>	Enters BGP configuration mode.
network <i>network mask mask</i>	Advertises a network into BGP.
show interfaces <i>interface</i>	Displays interface information and traffic statistics.
show tunnel 6rd	Displays information about 6RD tunnels.

Task 1: Configure a Static IPv6-in-IPv4 Tunnel

In this task, you will establish a static IPv6-in-IPv4 tunnel between two CE routers in different pods in the same team.

Activity Procedure

Complete these steps:

- Step 1** From the CE router, ping the other pod CE router using IPv4. Pings should be sourced from the Loopback 0 interface and destined to the Loopback 0 interface. The ping should be successful.

The CE router (Cisco IOS) output should be similar to the following, taken from Pod 3:

```
CE3# ping 10.4.10.1 source Loopback 0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.4.10.1, timeout is 2 seconds:
Packet sent with a source address of 10.3.10.1
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

The CE router (Cisco IOS) output should be similar to the following, taken from Pod 4:

```
CE4# ping 10.3.10.1 source Loopback 0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.3.10.1, timeout is 2 seconds:
Packet sent with a source address of 10.4.10.1
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

- Step 2** From the CE router, ping the other pod CE router using IPv6. Pings should be sourced from the Loopback 0 interface and destined to the Loopback 0 interface. The ping should be successful.
- The CE router (Cisco IOS) output should be similar to the following, taken from Pod 3:

```
CE3# ping 2001:DB8:10:4:10::1 source 2001:DB8:10:3:10::1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:10:4:10::1, timeout is 2 seconds:
Packet sent with a source address of 2001:DB8:10:3:10::1
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

The CE router (Cisco IOS) output should be similar to the following, taken from Pod 4:

```
CE4# ping 2001:DB8:10:3:10::1 source 2001:DB8:10:4:10::1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:10:3:10::1, timeout is 2 seconds:
Packet sent with a source address of 2001:DB8:10:4:10::1
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

- Step 3** On the CE router, create a tunnel 0 interface. Enable IPv6 on the interface. Use link-local IPv6 addresses for tunnel interface addressing.

Note	If you are unable to create the tunnel interface because of PIM registering, disable the PIM RP static IP address.
-------------	--

- Step 4** On the CE router, specify the Loopback 0 interface as the tunnel 0 source. Specify the IP address of the other pod CE router Loopback 0 interface as the tunnel destination.
- Step 5** On the CE router, set the tunnel 0 mode to IPv6-in-IPv4.
- Step 6** On the CE router, create a static IPv6 route for other pod CE router Loopback 0 interface that will point to the tunnel interface.
- Step 7** Before pinging, examine the tunnel interface traffic statistics. Observe the number of the packets that went into and out of the tunnel interface.

The CE router (Cisco IOS) output should be similar to the following, taken from Pod 3:

```
CE3# show interfaces Tunnel 0 | include put,
2785 packets input, 341640 bytes, 0 no buffer
3332 packets output, 338810 bytes, 0 underruns
```

The CE router (Cisco IOS) output should be similar to the following, taken from Pod 4:

```
CE4# show interfaces Tunnel 0 | include put,  
3340 packets input, 406718 bytes, 0 no buffer  
2830 packets output, 291080 bytes, 0 underruns
```

Step 8 From the CE router, ping the other pod CE router using IPv6. Pings should be sourced from the Loopback 0 interface and destined to the Loopback 0 interface. The ping should be successful.

NOTE: The other pod should be finished with the configuration of the CE router for the ping to be successful.

The CE router (Cisco IOS) output should be similar to the following, taken from Pod 3:

```
CE3# ping 2001:DB8:10:4:10::1 source 2001:DB8:10:3:10::1  
Type escape sequence to abort.  
Sending 5, 100-byte ICMP Echos to 2001:DB8:10:4:10::1, timeout is 2 seconds:  
Packet sent with a source address of 2001:DB8:10:3:10::1  
!!!!!  
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

The CE router (Cisco IOS) output should be similar to the following, taken from Pod 4:

```
CE4# ping 2001:DB8:10:3:10::1 source 2001:DB8:10:4:10::1  
Type escape sequence to abort.  
Sending 5, 100-byte ICMP Echos to 2001:DB8:10:3:10::1, timeout is 2 seconds:  
Packet sent with a source address of 2001:DB8:10:4:10::1  
!!!!!  
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

Step 9 Verify that pings went over the tunnel interface by examining the tunnel interface traffic statistics again. You should see that five additional packets went into and out of the tunnel interface.

The CE router (Cisco IOS) output should be similar to the following, taken from Pod 3:

```
CE3# show interfaces Tunnel 0 | include put,  
2790 packets input, 342340 bytes, 0 no buffer  
3337 packets output, 339410 bytes, 0 underruns
```

The CE router (Cisco IOS) output should be similar to the following, taken from Pod 4:

```
CE4# show interfaces Tunnel 0 | include put,  
3345 packets input, 407418 bytes, 0 no buffer  
2835 packets output, 291680 bytes, 0 underruns
```

Step 10 From the CE router, remove the tunnel interface and IPv6 static route that was created in this task.

Verification

No verification is needed for this task.

Task 2: Configure Dynamic 6RD Tunnels

In this task, you will deploy 6RD tunnels and configure a 6RD CE router and 6RD Border Relay. PE routers running Cisco IOS XE Software will act as 6RD Border Relay and CE routers will act as 6RD CE routers. You will configure 6RD tunnels between existing loopback interfaces.

Activity Procedure

Complete these steps:

Note	Coordinate this task with the other pod because only one PE router will act as a 6RD Border Relay. Remember that 6RD is not supported on routers running Cisco IOS XR Software.
-------------	---

Step 1 On the PE router (Cisco IOS XE), advertise the network configured on the Loopback 0 interface through BGP.

Step 2 From the CE router, ping the Loopback 0 interface of the PE (Cisco IOS XE) router. Use Loopback 0 as the source interface. The ping should be successful.

The CE router (Cisco IOS) output should be similar to the following, taken from Pod 3:

```
CE3# ping 10.4.1.1 source Loopback 0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.4.1.1, timeout is 2 seconds:
Packet sent with a source address of 10.3.10.1
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

The CE router (Cisco IOS) output should be similar to the following, taken from Pod 4:

```
CE4# ping 10.4.1.1 source Loopback 0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.4.1.1, timeout is 2 seconds:
Packet sent with a source address of 10.4.10.1
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

Step 3 How long is a prefix that is common to CEx, CEy, and PEy loopback interfaces?

Step 4 How many bits from the IPv4 address will be used to construct 6RD networks?

Step 5 Construct 6RD networks for IPv6 sites behind the CE routers. Use 2001:db8:aa00::/40 as the 6RD prefix and the last three octets of the IPv4 address on the loopback interface. Verify the results with the other pod in the team. Fill in the following table:

CE Router	IPv4 Tunnel Endpoint	6RD Prefix	6RD Network
CEx	10.x.10.1	2001:db8:aa00::/40	
CEy	10.y.10.1	2001:db8:aa00::/40	

Step 6 Also construct a 6RD network for the PE (Cisco IOS XE) router. This network will be used later to configure a default route on the CE routers. Verify the results with the other pod in the team.

PE Router	IPv4 Tunnel Endpoint	6RD Prefix	6RD Network
PEy	10.y.1.1	2001:db8:aa00::/40	

Step 7 On the CE router configure the tunnel interface with the following parameters:

- Enable the tunnel interface for IPv6
- **Tunnel source:** Loopback 0
- **Tunnel mode:** 6rd
- **Common prefix:** 8 bits
- **6RD prefix:** 2001:db8:aa00::/40
- **6RD BR:** 10.y.1.1

Step 8 On the CE router, verify information about the 6RD tunnel. Compare the displayed 6RD network with the network that you calculated in the previous steps.

The CE router (Cisco IOS) output should be similar to the following, taken from Pod 3:

```
CE3# show tunnel 6rd Tunnel0
Interface Tunnel0:
  Tunnel Source: 10.3.10.1
  6RD: Operational, V6 Prefix: 2001:DB8:AA00::/40
    V4 Prefix, Length: 8, Value: 10.0.0.0
    V4 Suffix, Length: 0, Value: 0.0.0.0
    Border Relay address: 10.4.1.1
  General Prefix: 2001:DB8:AA03:A01::/64
```

The CE router (Cisco IOS) output should be similar to the following, taken from Pod 4:

```
CE4# show tunnel 6rd Tunnel0
Interface Tunnel0:
  Tunnel Source: 10.4.10.1
  6RD: Operational, V6 Prefix: 2001:DB8:AA00::/40
    V4 Prefix, Length: 8, Value: 10.0.0.0
    V4 Suffix, Length: 0, Value: 0.0.0.0
    Border Relay address: 10.4.1.1
  General Prefix: 2001:DB8:AA04:A01::/64
```

- Step 9** On the CE router, assign the first IPv6 address from the 6RD network to a new loopback interface. Use 10 as the interface identifier.
- Step 10** On the CE router, create an IPv6 static route for the 6RD prefix that will use the tunnel 0 interface as the outgoing interface.
- Step 11** On the PE (Cisco IOS XE) router, configure the tunnel 10 interface with the following parameters:
- Enable the tunnel interface for IPv6
 - **Tunnel source:** Loopback 0
 - **Tunnel mode:** 6rd
 - **Common prefix:** 8 bits
 - **6RD prefix:** 2001:db8:aa00::/40
- Step 12** On the PE router create an IPv6 static route for the 6RD prefix that will use the tunnel 10 interface as the outgoing interface.
- Step 13** On the PE (Cisco IOS XE) router, verify the information about the 6RD tunnel. Compare the displayed 6RD network with the network that you calculated in the previous steps.
- The PE router (Cisco IOS XE) output should be similar to the following, taken from Pod 4:

```
PE4# show tunnel 6rd
Interface Tunnel10:
  Tunnel Source: 10.4.1.1
  6RD: Operational, V6 Prefix: 2001:DB8:AA00::/40
    V4 Prefix, Length: 8, Value: 10.0.0.0
    V4 Suffix, Length: 0, Value: 0.0.0.0
  General Prefix: 2001:DB8:AA04:101::/64
```

- Step 14** On the PE (Cisco IOS XE) router, create another loopback interface. Use 10 as the interface identifier and 2001:db8:101:y::1 as the IP address on the interface. This interface will present the IPv6 Internet that is available over the 6RD Border Relay.
- Step 15** On the CE router, create a default route that will use the tunnel interface as the outgoing interface and will point to the 6RD network of the PE (Cisco IOS XE) router.
- Step 16** Before pinging, examine the tunnel interface traffic statistics. Observe the number of the packets that went into and out of the tunnel interface.

The CE router (Cisco IOS) output should be similar to the following, taken from Pod 3:

```
CE3# show interfaces Tunnel0 | include packets
5 minute input rate 1000 bits/sec, 1 packets/sec
5 minute output rate 0 bits/sec, 0 packets/sec
20 packets input, 2800 bytes, 0 no buffer
168 packets output, 17234 bytes, 0 underruns
```

The CE router (Cisco IOS) output should be similar to the following, taken from Pod 4:

```
CE4# show interfaces Tunnel0 | include packets
5 minute input rate 1000 bits/sec, 1 packets/sec
5 minute output rate 0 bits/sec, 0 packets/sec
30 packets input, 4200 bytes, 0 no buffer
170 packets output, 17430 bytes, 0 underruns
```

Step 17 From the CE router, ping the other pod CE router Loopback10 interface. Use Loopback10 as a source interface. The ping should be successful and traffic should go directly to the other pod CE router.

Note: The other pod should be finished with the configuration of the CE router for the ping to be successful.

The CE router (Cisco IOS) output should be similar to the following, taken from Pod 3:

```
CE3# ping 2001:DB8:AA04:A01::1 source Loopback 10
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:AA04:A01::1, timeout is 2 seconds:
Packet sent with a source address of 2001:DB8:AA03:A01::1
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

The CE router (Cisco IOS) output should be similar to the following, taken from Pod 4:

```
CE4# ping 2001:DB8:AA03:A01::1 source Loopback 10
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:AA03:A01::1, timeout is 2 seconds:
Packet sent with a source address of 2001:DB8:AA04:A01::1
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

Step 18 Verify that pings went over the tunnel interface by examining the tunnel interface traffic statistics again. You should see that five additional packets went into and out of the tunnel interface.

The CE router (Cisco IOS) output should be similar to the following, taken from Pod 3:

```
CE3# show interfaces Tunnel0 | include packets
5 minute input rate 0 bits/sec, 0 packets/sec
5 minute output rate 0 bits/sec, 0 packets/sec
25 packets input, 3500 bytes, 0 no buffer
173 packets output, 17834 bytes, 0 underruns
```

The CE router (Cisco IOS) output should be similar to the following, taken from Pod 4:

```
CE4# show interfaces Tunnel0 | include packets
5 minute input rate 1000 bits/sec, 1 packets/sec
5 minute output rate 0 bits/sec, 0 packets/sec
35 packets input, 4900 bytes, 0 no buffer
175 packets output, 18030 bytes, 0 underruns
```

Step 19 Before pinging a second time, examine the tunnel interface traffic statistics. Observe the number of the packets that went into and out of the tunnel interface.

The CE router (Cisco IOS) output should be similar to the following, taken from Pod 3:

```
CE3# show interfaces Tunnel0 | include packets
 5 minute input rate 0 bits/sec, 0 packets/sec
 5 minute output rate 0 bits/sec, 0 packets/sec
 50 packets input, 7000 bytes, 0 no buffer
 361 packets output, 36752 bytes, 0 underruns
```

The CE router (Cisco IOS) output should be similar to the following, taken from Pod 4:

```
CE4# show interfaces Tunnel0 | include packets
 5 minute input rate 0 bits/sec, 0 packets/sec
 5 minute output rate 0 bits/sec, 0 packets/sec
 45 packets input, 6300 bytes, 0 no buffer
 342 packets output, 34616 bytes, 0 underruns
```

Step 20 From the CE router, ping the other pod PE (Cisco IOS XE) router Loopback10 interface. Use Loopback10 as a source interface. The ping should be successful.

Note: The other pod should be finished with the configuration of the CE router for the ping to be successful.

The CE router (Cisco IOS) output should be similar to the following, taken from Pod 3:

```
CE3# ping 2001:DB8:101:4::1 source Loopback 10
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:101:4::1, timeout is 2 seconds:
Packet sent with a source address of 2001:DB8:AA03:A01::1
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

The CE router (Cisco IOS) output should be similar to the following, taken from Pod 4:

```
CE4# ping 2001:DB8:101:4::1 source Loopback 10
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:101:4::1, timeout is 2 seconds:
Packet sent with a source address of 2001:DB8:AA04:A01::1
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

Step 21 Verify that pings went over the tunnel interface by examining the tunnel interface traffic statistics again. You should see that five additional packets went into and out of the tunnel interface.

The CE router (Cisco IOS) output should be similar to the following, taken from Pod 3:

```
CE3# show interfaces Tunnel0 | include packets
 5 minute input rate 0 bits/sec, 0 packets/sec
 5 minute output rate 0 bits/sec, 0 packets/sec
 55 packets input, 7700 bytes, 0 no buffer
 366 packets output, 37352 bytes, 0 underruns
```

The CE router (Cisco IOS) output should be similar to the following, taken from Pod 4:

```
CE4# show interfaces Tunnel0 | include packets
5 minute input rate 1000 bits/sec, 1 packets/sec
5 minute output rate 0 bits/sec, 0 packets/sec
50 packets input, 7000 bytes, 0 no buffer
347 packets output, 35216 bytes, 0 underruns
```

Verification

No verification is needed for this task.

Answer Key

Hardware Lab 1: Implement BGP Route Reflectors

Task 2: Configure a Route Reflector and Internal BGP Session

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

```
RP/0/RSP0/CPU0:PEx(config)#
interface GigabitEthernet0/0/0/3
  ipv4 address 192.168.x2.x0 255.255.255.0
  ipv6 address 2001:db8:192:168:x2::x0/80
  no shutdown
!
router isis 1
  interface GigabitEthernet0/0/0/3
    circuit-type level-2-only
    address-family ipv4 unicast
    address-family ipv6 unicast
!
commit
```

```
PEy(config)#
interface GigabitEthernet0/0/3
  ip address 192.168.y2.y0 255.255.255.0
  ipv6 enable
  ipv6 address 2001:db8:192:168:y2::y0/80
  ip router isis
  ipv6 router isis
  isis circuit-type level-2-only
```

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

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

```
RP/0/RSP0/CPU0:PEx(config)#
router bgp 64500
 neighbor 10.0.2.1
   remote-as 64500
   update-source Loopback0
 address-family ipv4 unicast
!
commit
```

```
PEy(config)#
router bgp 64500
 neighbor 10.0.2.1 remote-as 64500
 neighbor 10.0.2.1 update-source Loopback0
```

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

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

```
RP/0/RSP0/CPU0:PEx(config)#
router bgp 64500
 neighbor 10.0.2.1
   address-family ipv4 unicast
   next-hop-self
!
commit
```

```
PEy(config)#
router bgp 64500
 neighbor 10.0.2.1 next-hop-self
```

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

Step 4 Enter these commands on the P2 router, running IOS XR.

```
RP/0/RSP0/CPU0:P2(config)#
router bgp 64500
 bgp cluster-id 10.0.1.1
 neighbor 10.x.1.1
   remote-as 64500
   update-source Loopback0
   address-family ipv4 unicast
   route-reflector-client
!
commit
```

or

```
RP/0/RSP0/CPU0:P2(config)#
router bgp 64500
  bgp cluster-id 10.0.1.1
  neighbor 10.y.1.1
    remote-as 64500
    update-source Loopback0
    address-family ipv4 unicast
      route-reflector-client
    !
  commit
```

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

Step 8 The lowest router ID among the sending routers.

Task 3: Restrict Route Propagation to a Client

Step 1 Enter these commands on the P2 router, running IOS XR.

```
RP/0/RSP0/CPU0:P2(config)#
route-policy FILTER_TO_CLIENT_PODx
  if as-path originates-from '6450y' then
    pass
  endif
end-policy
!
commit
```

or

```
RP/0/RSP0/CPU0:P2(config)#
route-policy FILTER_TO_CLIENT_PODy
  if as-path originates-from '6450x' then
    pass
  endif
end-policy
!
commit
```

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

Step 2 Enter these commands on the P2 router, running IOS XR.

```
RP/0/RSP0/CPU0:P2(config)#
router bgp 64500
  neighbor 10.x.1.1
    address-family ipv4 unicast
      route-policy FILTER_TO_CLIENT_PODx out
    !
  commit
```

or

```
RP/0/RSP0/CPU0:P2(config)#
router bgp 64500
  neighbor 10.y.1.1
    address-family ipv4 unicast
      route-policy FILTER_TO_CLIENT_PODy out
!
commit
```

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

Clear BGP sessions on the P2 router.

```
RP/0/RSP0/CPU0:P2# clear ip bgp *
Fri May 23 04:54:57.475 UTC
```

Verification of this task is possible only when at least one pod from the other team has a working BGP configuration.

Hardware Lab 2: Implement BGP Security Options

Task 1: Implement BGP Neighbor Authentication Using Passwords

Step 2 Enter these commands on the CE routers, running Cisco IOS Software:

```
CEx(config)#
router bgp 6450x
  neighbor 192.168.10x.x0 password C!sc()
```

```
CEy(config)#
router bgp 6450y
  neighbor 192.168.10y.y0 password C!sc()
```

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

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

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

```
PEy(config)#
router bgp 64500
  neighbor 192.168.10y.y1 password C!sc()
```

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

Task 2: Implement BGP Neighbor Authentication Using Key Chains

Step 2 Enter these commands on the PE routers, running IOS XR:

```
RP/0/RSP0/CPU0:PEx(config)#
key chain BGP
  key 1
    accept-lifetime 00:00:00 january 01 2014 infinite
    key-string C!sc()
    send-lifetime 00:00:00 january 01 2014 infinite
    cryptographic-algorithm HMAC-MD5
  !
commit
```

x should be substituted with 1 or 3, depending on your team assignment.

NOTE: Because Cisco IOS XR Software encrypts all passwords by default, the **show key chain BGP** command would show the encrypted password instead of clear-text "C!sc()".

Step 3 Enter these commands on the PE routers, running IOS XR:

```
RP/0/RSP0/CPU0:PEx(config)#
router bgp 64500
  neighbor 10.0.1.1
    keychain BGP
  !
commit
```

x should be substituted with 1 or 3, depending on your team assignment.

Step 4 Enter these commands on the P1 router, running IOS XR:

```
RP/0/RSP0/CPU0:P1(config)#
key chain BGP_PODx
  key 1
    accept-lifetime 00:00:00 january 01 2014 infinite
    key-string C!sc()
    send-lifetime 00:00:00 january 01 2014 infinite
    cryptographic-algorithm HMAC-MD5
  !
commit
```

x should be substituted with 1 or 3, depending on your team assignment.

Step 5 Enter these commands on the P1 router, running IOS XR:

```
RP/0/RSP0/CPU0:P1(config)#
router bgp 64500
  neighbor 10.x.1.1
  keychain BGP_PODx
!
commit
```

x should be substituted with 1 or 3, depending on your team assignment.

Task 3: Enable BGP TTL Security Check

Step 1 Enter these commands on the CE routers, running IOS:

```
CEx(config)#
router bgp 6450x
  neighbor 192.168.10x.x0 ttl-security hops 1
```

```
CEy(config)#
router bgp 6450y
  neighbor 192.168.10y.y0 ttl-security hops 1
```

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

Step 2 1, because the neighbors are directly connected.

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

```
RP/0/RSP0/CPU0:PEx(config)#
router bgp 64500
  neighbor 192.168.10x.x1
  ttl-security
!
commit
```

```
PEy(config)#
router bgp 64500
  neighbor 192.168.10y.y1 ttl-security hops 1
```

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

Task 4: Enable CoPP

Step 1 Enter these commands on the CE routers, running Cisco IOS Software:

```
CEx(config)#  
ip access-list extended BGP_TRAFFIC  
  permit tcp host 192.168.10x.x0 host 192.168.10x.x1 eq bgp  
  permit tcp host 192.168.10x.x0 eq bgp host 192.168.10x.x1
```

```
CEy(config)#  
ip access-list extended BGP_TRAFFIC  
  permit tcp host 192.168.10y.y0 host 192.168.10y.y1 eq bgp  
  permit tcp host 192.168.10y.y0 eq bgp host 192.168.10y.y1
```

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

Step 2 Enter these commands on the CE routers, running Cisco IOS Software:

```
CEx(config)#  
class-map match-all BGP_CLASS  
  match access-group name BGP_TRAFFIC
```

```
CEy(config)#  
class-map match-all BGP_CLASS  
  match access-group name BGP_TRAFFIC
```

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

Step 3 Enter these commands on the CE routers, running Cisco IOS Software:

```
CEx(config)#  
policy-map COPP  
  class BGP_CLASS  
    police rate 200 pps conform-action transmit exceed-action drop
```

```
CEy(config)#  
policy-map COPP  
  class BGP_CLASS  
    police rate 200 pps conform-action transmit exceed-action drop
```

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

Step 4 Enter these commands on the CE routers, running Cisco IOS Software:

```
CEx(config)#  
control-plane  
  service-policy input COPP
```

```
CEy(config)#  
control-plane  
service-policy input COPP
```

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

Task 5: Enable RTBH Filtering

Step 1 Enter these commands on the CE routers, running Cisco IOS Software:

```
CEx(config)#  
interface Loopback1  
ip address 10.x.100.1 255.255.255.255
```

```
CEy(config)#  
interface Loopback1  
ip address 10.y.100.1 255.255.255.255
```

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

Step 2 Enter these commands on the CE routers, running Cisco IOS Software:

```
CEx(config)#  
router bgp 6450x  
address-family ipv4  
network 10.x.100.1 mask 255.255.255.255
```

```
CEy(config)#  
router bgp 6450y  
address-family ipv4  
network 10.y.100.1 mask 255.255.255.255
```

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

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

```
RP/0/RSP0/CPU0:PEx(config)#  
router static  
address-family ipv4 unicast  
172.16.x.0/24 Null0  
!  
commit
```

```
PEy(config)#  
ip route 172.16.y.0 255.255.255.0 Null 0
```

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

Step 5 Enter these commands on the PE routers, running IOS XR and IOS XE, respectively.

```
RP/0/RSP0/CPU0:PEx(config)#  
interface GigabitEthernet0/0/0/0  
  ipv4 verify unicast source reachable-via rx  
!  
commit
```

```
PEy(config)#  
interface GigabitEthernet0/0/0  
  ip verify unicast source reachable-via rx
```

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

Step 6 Enter these commands on the P1 router, running IOS XR:

```
RP/0/RSP0/CPU0:P1(config)#  
router static  
  address-family ipv4 unicast  
    172.16.x.0/24 Null0  
!  
commit
```

```
RP/0/RSP0/CPU0:P1(config)#  
router static  
  address-family ipv4 unicast  
    172.16.y.0/24 Null0  
!  
commit
```

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

Step 9 The black-holed routes should not be advertised outside the autonomous system, therefore they are tagged with no-export community.

Step 10 To prefer an already existing route for the same network that has been reflected from the R2 route reflector and points to the CE router on the PE router.

Step 11 Enter these commands on the P1 router, running IOS XR:

```
RP/0/RSP0/CPU0:P1(config)#  
router static  
  address-family ipv4 unicast  
    10.x.100.1/32 Null0 tag 5  
!  
commit
```

```
RP/0/RSP0/CPU0:Pl(config)#  
router static  
  address-family ipv4 unicast  
    10.y.100.1/32 Null0 tag 5  
!  
commit
```

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

Hardware Lab 3: Improve BGP Scalability

Task 1: Implement the BGP Configuration and Peer Templates

Step 3 Enter these commands on the PE router, running IOS XR:

```
RP/0/RSP0/CPU0:PEX(config)#
router bgp 64500
  af-group IPv4 address-family ipv4 unicast
  route-policy pass in
  route-policy pass out
  !
  neighbor-group EBGp
  password C!sc()
  ttl-security
  address-family ipv4 unicast
  use af-group IPv4
  !
  no neighbor 192.168.10x.x1
  neighbor 192.168.10x.x1
  remote-as 6450x
  use neighbor-group EBGp
  !
commit
```

x should be substituted with 1 or 3, depending on your team assignment.

Step 4 Enter these commands on the PE router, running IOS XE:

```
PEY(config)#
router bgp 64500
  template peer-policy EBGp_POLICY
  exit-peer-policy
  !
  template peer-session EBGp_SESSION
  password C!sc()
  ttl-security hops 1
  exit-peer-session
  !
  no neighbor 192.168.10y.y1 remote-as 6450y
  neighbor 192.168.10y.y1 remote-as 6450y
  neighbor 192.168.10y.y1 inherit peer-session EBGp_SESSION
  !
  address-family ipv4
  neighbor 192.168.10y.y1 inherit peer-policy EBGp_POLICY
  exit-address-family
```

y should be substituted with 2 or 4, depending on your team assignment.

Task 2: Limit the Number of Prefixes Received from a BGP Neighbor

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

```
RP/0/RSP0/CPU0:PEx(config)#
router bgp 64500
  af-group IPv4 address-family ipv4 unicast
    maximum-prefix 2
!
commit
```

```
PEy(config)#
router bgp 64500
  template peer-policy EBG_POLICY
    maximum-prefix 2
```

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

NOTE: When you set BGP maximum prefix 2, the output in the running configuration shows the **maximum-prefix 2 75** command. The value 75 is the default, which means that the 75% threshold is taken to generate a warning message.

Task 3: Improve BGP Convergence by Changing the BGP Scan and Advertisement Interval

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

```
RP/0/RSP0/CPU0:PEx(config)#
router bgp 64500
  bgp scan-time 30
!
commit
```

```
PEy(config)#
router bgp 64500
  bgp scan-time 30
```

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

Step 5 Enter these commands on the CE routers, running IOS.

```
CEx(config)#
router bgp 6450x
  neighbor 192.168.10x.x0 advertisement-interval 15
```

```
CEy(config)#  
router bgp 6450y  
neighbor 192.168.10y.y0 advertisement-interval 15
```

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

Step 7 To disable the Loopback0 interface, enter these commands on the CE routers, running IOS.

```
CEx(config)#  
interface Loopback 0  
shutdown
```

```
CEy(config)#  
interface Loopback 0  
shutdown
```

To enable the Loopback0 interface, enter these commands on the CE routers, running IOS.

```
CEx(config)#  
interface Loopback 0  
no shutdown
```

```
CEy(config)#  
interface Loopback 0  
no shutdown
```

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

Task 4: Implement BGP Route Dampening

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

```
RP/0/RSP0/CPU0:PEx(config)#  
router bgp 64500  
address-family ipv4 unicast  
bgp dampening  
!  
commit  
!  
end  
!  
debug bgp dampening
```

```
PEy(config)#
router bgp 64500
  address-family ipv4
    bgp dampening
  end
!
```

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

Step 2 Enter these commands on the CE routers, running IOS.

```
CEx(config)#
interface Loopback0
shutdown
```

```
CEy(config)#
interface Loopback0
shutdown
```

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

Step 6 The route needs to flap three times before it is suppressed.

Step 7 The default suppress penalty is 2000. The half-life time is 15 minutes. The reuse penalty is 750.

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

```
RP/0/RSP0/CPU0 :PEx#
undebug all
```

```
PEy#
undebug all
```

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

Hardware Lab 4: Implement Layer 2 and Layer 3 Multicast

Task 1: Enable IGMP and MLD

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

```
RP/0/RSP0/CPU0: PEx(config)#
multicast-routing
 address-family ipv4
   interface GigabitEthernet0/0/0/0
     enable
 !
 address-family ipv6
   interface GigabitEthernet0/0/0/0
     enable
 !
commit
```

```
PEy(config)#
ip multicast-routing distributed
ipv6 multicast-routing
```

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

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

```
RP/0/RSP0/CPU0: PEx(config)#
router igmp
 interface GigabitEthernet0/0/0/0
   version 2
   query-interval 30
 !
commit
```

```
PEy(config)#
interface GigabitEthernet0/0/0
ip igmp version 2
ip igmp query-interval 30
```

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

Step 5 Enter these commands on the PE routers, running IOS XR and IOS XE, respectively.

```
RP/0/RSP0/CPU0: PEx(config)#
router mld
 interface GigabitEthernet0/0/0/0
   query-interval 60
   router enable
 !
commit
```

```
PEy(config)#  
interface GigabitEthernet0/0/0  
  ipv6 mld router  
  ipv6 mld query-interval 60
```

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

Step 6 Enter these commands on the CE routers, running Cisco IOS.

```
CEx(config)#  
interface GigabitEthernet0/0  
  ip igmp join-group 234.1.1.1
```

```
CEy(config)#  
interface GigabitEthernet0/0  
  ip igmp join-group 234.1.1.1
```

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

Task 2: Verify IGMP Snooping

Step 2 Enter these commands on the SW switch:

```
SWx(config)#  
no ip igmp snooping
```

```
SWy(config)#  
no ip igmp snooping
```

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

Step 4 Enter these commands on the SW switch:

```
SWx(config)#  
ip igmp snooping
```

```
SWy(config)#  
ip igmp snooping
```

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

Step 5 Enter these commands on the CE routers, running Cisco IOS.

```
CEx(config)#  
interface GigabitEthernet0/0  
no ip igmp join-group 234.1.1.1
```

```
CEy(config)#  
interface GigabitEthernet0/0  
no ip igmp join-group 234.1.1.1
```

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

Hardware Lab 5: Enable and Optimize PIM-SM

Task 1: Implement PIM-SM

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

```
CEx(config)#  
ip multicast-routing
```

```
RP/0/RSP0/CPU0:PEx(config)#  
multicast-routing  
address-family ipv4  
interface Loopback0  
enable  
!  
interface GigabitEthernet0/0/0/0  
enable  
!  
interface GigabitEthernet0/0/0/1  
enable  
!  
interface GigabitEthernet0/0/0/2  
enable  
!  
interface GigabitEthernet0/0/0/3  
enable  
!  
commit
```

```
CEy(config)#  
ip multicast-routing
```

```
PEy(config)#  
ip multicast-routing distributed
```

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

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

```
CEx(config)#  
interface Loopback0  
ip pim sparse-mode  
!  
interface GigabitEthernet0/0  
ip pim sparse-mode
```

```

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

```

```

CEy(config)#
interface Loopback0
  ip pim sparse-mode
!
interface GigabitEthernet0/0
  ip pim sparse-mode

```

```

PEy(config)#
interface Loopback0
  ip pim sparse-mode
!
interface GigabitEthernet0/0/0
  ip pim sparse-mode
!
interface GigabitEthernet0/0/1
  ip pim sparse-mode
!
interface GigabitEthernet0/0/2
  ip pim sparse-mode
!
interface GigabitEthernet0/0/3
  ip pim sparse-mode

```

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

- Step 5** Enter these commands on the CE and PE routers, running IOS, IOS XR, and IOS XE, respectively.

```

CEx(config)#
ip pim spt-threshold infinity

```

```
RP/0/RSP0/CPU0:PEx(config)#
router pim
  address-family ipv4
    spt-threshold infinity
!
commit
```

```
CEy(config)#
ip pim spt-threshold infinity
```

```
PEy(config)#
ip pim spt-threshold infinity
```

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

- Step 6** Enter these commands on the CE and PE routers, running IOS, IOS XR, and IOS XE, respectively.

```
CEx(config)#
ip pim rp-address 10.0.1.1
```

```
RP/0/RSP0/CPU0:PEx(config)#
router pim
  address-family ipv4
    rp-address 10.0.1.1
!
commit
```

```
CEy(config)#
ip pim rp-address 10.0.1.1
```

```
PEy(config)#
ip pim rp-address 10.0.1.1
```

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

Task 2: Shared Tree Formation—Receivers

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

```
CEx(config)#
interface Loopback0
 ip igmp join-group 224.1.x.1
```

```
RP/0/RSP0/CPU0:PEx(config)#
router igmp
 interface Loopback0
   join-group 224.1.x.1
!
commit
```

```
CEy(config)#
interface Loopback0
 ip igmp join-group 224.1.y.1
```

```
PEy(config)#
interface Loopback0
 ip igmp join-group 224.1.y.1
```

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

Step 3 The results for group 224.1.3.1 from the CE, PE, and P1 routers should be similar to the following, taken from Pod 3:

Router	Incoming Interface	OIL
CE3	GigabitEthernet0/0	Loopback0
PE3	GigabitEthernet0/0/0/2	Loopback0, GigabitEthernet0/0/0/0
P1	Decapstunnel0	GigabitEthernet0/0/0/3

The results for group 224.1.4.1 from CE, PE, and P1 routers should be similar to the following, taken from Pod 4:

Router	Incoming Interface	OIL
CE4	GigabitEthernet0/0	Loopback0
PE4	GigabitEthernet0/0/2	Loopback0, GigabitEthernet0/0/0/0
P1	Decapstunnel0	GigabitEthernet0/0/0/5

Step 4 P1 acts as the RP. Because no sources are active, the RP does not receive any traffic for the group. Therefore, there are no incoming interfaces for the group.

Task 3: Shared Tree Formation—Sources

Step 1 Enter these commands on the CE routers, running IOS.

```
CEx(config)#  
router bgp 6450x  
address-family ipv4  
network 192.168.10x.0 mask 255.255.255.0
```

```
CEy(config)#  
router bgp 6450y  
address-family ipv4  
network 192.168.10y.0 mask 255.255.255.0
```

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

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

```
RP/0/RSP0/CPU0:PEx(config)#  
router bgp 64500  
af-group IPv4 address-family ipv4 unicast  
no maximum-prefix 2 75  
!  
commit
```

```
PEy(config)#  
router bgp 64500  
template peer-policy EBGP_POLICY  
no maximum-prefix 2
```

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

- Step 6** The entry did not change. The CE router uses shared tree to receive multicast traffic, and the active source does not influence the mroute table on the router.
- Step 7** No (S,G) are entries present. Because the CE router uses the shared tree, the only entry for the group is (*,G).
- Step 9** GigabitEthernet0/0/0
- Step 10** GigabitEthernet0/0/2
- Step 11** The PE router would use the (*,G) entry to receive multicast traffic from the RP. Because no receivers are present in the path to the RP through the PE router, the OIL list is empty.

Task 4: Switching to the SPT

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

```
CEx(config)#
ip pim spt-threshold 0
```

```
RP/0/RSP0/CPU0: PEx(config)#
router pim
address-family ipv4
no spt-threshold infinity
!
commit
```

```
CEy(config)#
ip pim spt-threshold 0
```

```
PEy(config)#
ip pim spt-threshold 0
```

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

- Step 4** Yes. Because the CE router created the SPT, the (S,G) entry exists.
- Step 6** Yes. Because the PE router created the SPT, the (S,G) entry exists.
- Step 7** GigabitEthernet0/0/0/1
- You should see that traffic now flows between the PE routers directly because the SPT has been built. When the SPT switchover was disabled, traffic went over the RP router.
- Step 8** Enter these commands on the CE and PE routers, running IOS, IOS XR, and IOS XE, respectively.

```
CEx(config)#
interface Loopback0
no ip igmp join-group 224.1.x.1
```

```
RP/0/RSP0/CPU0:PEx(config)#  
router igmp  
  interface Loopback0  
    no join-group 224.1.x.1  
!  
commit
```

```
CEy(config)#  
interface Loopback0  
  no ip igmp join-group 224.1.y.1
```

```
PEy(config)#  
interface Loopback0  
  no ip igmp join-group 224.1.y.1
```

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

Hardware Lab 6: Implement PIM-SM Enhancements

Task 1: Implement PIM-SSM

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

```
CEx(config)#
ip access-list standard SSM_PODx
  permit 224.1.x.1
  permit 224.1.y.1
!
ip pim ssm range SSM_PODx
```

```
RP/0/RSP0/CPU0:PEx(config)#
ipv4 access-list SSM_PODx
  10 permit 224.1.x.1
  20 permit 224.1.y.1
!
multicast-routing
  ipv4 ssm range SSM_PODx
!
commit
```

```
CEy(config)#
ip access-list standard SSM_PODy
  permit 224.1.x.1
  permit 224.1.y.1
!
ip pim ssm range SSM_PODy
```

```
PEy(config)#
ip access-list standard SSM_PODy
  permit 224.1.x.1
  permit 224.1.y.1
!
ip pim ssm range SSM_PODy
```

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

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

```
CEx(config)#
interface Loopback0
  ip igmp join-group 224.1.x.1 source 192.168.10y.y1
```

```
RP/0/RSP0/CPU0:PEx(config)#
router igmp
 interface Loopback0
  join-group 224.1.x.1 192.168.10y.y1
!
commit
```

```
CEy(config)#
interface Loopback0
 ip igmp join-group 224.1.y.1 source 192.168.10x.x1
```

```
PEy(config)#
interface Loopback0
 ip igmp join-group 224.1.y.1 source 192.168.10x.x1
```

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

Step 6 GigabitEthernet0/0/0/1

Step 7 GigabitEthernet0/0/0/0, Loopback0

Step 9 GigabitEthernet0/0/0

Step 10 GigabitEthernet0/0/1

Step 12 Enter these commands on the CE and PE routers, running IOS, IOS XR, and IOS XE, respectively.

```
CEx(config)#
interface Loopback0
 no ip igmp join-group 224.1.x.1 source 192.168.10y.y1
```

```
RP/0/RSP0/CPU0:PEx(config)#
router igmp
 interface Loopback0
  no join-group 224.1.x.1 192.168.10y.y1
!
commit
```

```
CEy(config)#
interface Loopback0
 no ip igmp join-group 224.1.y.1 source 192.168.10x.x1
```

```
PEy(config)#
interface Loopback0
 no ip igmp join-group 224.1.y.1 source 192.168.10x.x1
```

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

Step 13 Enter these commands on the CE and PE routers, running IOS, IOS XR, and IOS XE, respectively.

```
CEx(config)#  
no ip pim ssm range SSM_PODx  
!  
no ip access-list standard SSM_PODx
```

```
RP/0/RSP0/CPU0:PEx(config)#  
multicast-routing  
  address-family ipv4  
    no ssm range SSM_PODx  
  !  
no ipv4 access-list SSM_PODx  
!  
commit
```

```
CEy(config)#  
no ip pim ssm range SSM_PODy  
!  
no ip access-list standard SSM_PODy
```

```
PEy(config)#  
no ip pim ssm range SSM_PODy  
!  
no ip access-list standard SSM_PODy
```

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

Hardware Lab 7: Implement Rendezvous Point Distribution

Task 1: Enable OSPF Routing

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

```
RP/0/RSP0/CPU0:PEx(config)#  
interface GigabitEthernet0/0/0/2  
shutdown  
!  
interface GigabitEthernet0/0/0/3  
shutdown  
!  
commit
```

```
PEy(config)#  
interface GigabitEthernet0/0/2  
shutdown  
!  
interface GigabitEthernet0/0/3  
shutdown
```

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

Step 2 Enter these commands on the CE router, running IOS.

```
CEx(config)#  
router ospf 1  
network 10.x.10.1 0.0.0.0 area 0  
network 192.168.10x.x1 0.0.0.0 area 0
```

```
CEy(config)#  
router ospf 1  
network 10.y.10.1 0.0.0.0 area 0  
network 192.168.10y.y1 0.0.0.0 area 0
```

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

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

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

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

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

Step 4 Enter these commands on the CE router, running IOS.

```
CEx(config)#
router bgp 6450x
neighbor 192.168.10x.x0 shutdown
neighbor 2001:DB8:192:168:10x::x0 shutdown
```

```
CEy(config)#
router bgp 6450y
neighbor 192.168.10y.y0 shutdown
neighbor 2001:DB8:192:168:10y::y0 shutdown
```

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

Step 6 Enter these commands on the CE and PE routers, running IOS, IOS XR, and IOS XE, respectively.

```
CEx(config)#
no ip pim rp-address 10.0.x.1
```

```
RP/0/RSP0/CPU0:PEx(config)#
router pim
address-family ipv4
no rp-address 10.0.1.1
!
commit
```

```
CEy(config)#  
no ip pim rp-address 10.0.y.1
```

```
PEy(config)#  
no ip pim rp-address 10.0.1.1
```

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

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

```
CEx(config)#  
interface Loopback0  
ip igmp join-group 224.1.x.1 source 192.168.10y.y1
```

```
RP/0/RSP0/CPU0:PEx(config)#  
router igmp  
interface Loopback0  
join-group 224.1.x.1 192.168.10y.y1  
!  
commit
```

```
CEy(config)#  
interface Loopback0  
ip igmp join-group 224.1.y.1 source 192.168.10x.x1
```

```
PEy(config)#  
interface Loopback0  
ip igmp join-group 224.1.y.1 source 192.168.10x.x1
```

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

Task 2: Enable BSR

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

```
RP/0/RSP0/CPU0:PEx(config)#  
router pim  
address-family ipv4  
bsr candidate-bsr 10.x.1.1  
bsr candidate-rp 10.x.1.1  
!  
commit
```

```
PEy(config)#  
ip pim bsr-candidate Loopback 0  
ip pim rp-candidate Loopback 0
```

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

NOTE: When the BSR candidate is configured, you can specify additional parameters such as the hash mask length and priority; otherwise the default values are used (length30 for the hash mask length and 1 for priority). Also, when the BSR RP candidate is configured, you can specify additional parameters such as priority and interval; otherwise the default values are used (-192 for priority and 60 for the interval).

You can issue the **show running-configuration router pim** command to verify configured parameters.

Task 3: Enable Anycast RP

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

```
RP/0/RSP0/CPU0:PEx(config)#  
router pim  
  address-family ipv4  
    no bsr candidate-bsr 10.x.1.1  
    no bsr candidate-rp 10.x.1.1  
  !  
commit
```

```
PEy(config)#  
no ip pim bsr-candidate Loopback 0  
no ip pim rp-candidate Loopback 0
```

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

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

```
RP/0/RSP0/CPU0:PEx(config)#  
interface Loopback1  
  ipv4 address 1.1.1.1 255.255.255.255  
  !  
multicast-routing  
  address-family ipv4  
    interface Loopback1  
      enable  
  !  
router ospf 1  
  area 0  
    interface Loopback1  
  !  
router pim  
  address-family ipv4  
    interface Loopback1  
      enable  
  !  
commit
```

```

PEy(config)#
interface Loopback1
 ip address 1.1.1.1 255.255.255.255
 ip pim sparse-mode
!
router ospf 1
 network 1.1.1.1 0.0.0.0 area 0

```

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

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

```

CEx(config)#
ip pim rp-address 1.1.1.1

```

```

RP/0/RSP0/CPU0:PEx(config)#
router pim
 address-family ipv4
  rp-address 1.1.1.1
!
commit

```

```

CEy(config)#
ip pim rp-address 1.1.1.1

```

```

PEy(config)#
ip pim rp-address 1.1.1.1

```

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

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

```

RP/0/RSP0/CPU0:PEx(config)#
router msdp
 originator-id Loopback0
 peer 10.y.1.1
  connect-source Loopback0
!
commit

```

```

PEy(config)#
ip msdp peer 10.x.1.1 connect-source Loopback0
ip msdp originator-id Loopback0

```

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

Hardware Lab 8: Implement a DHCPv6 Server with Prefix Delegation

Task 1: Configure a Prefix Delegation DHCPv6 Server and Client

Step 2 Enter these commands on the PE router, running IOS XE:

```
PEy(config)#
ipv6 local pool PREFIX 2001:DB8:100:y::/64 64
!
ipv6 dhcp pool Customers
prefix-delegation pool PREFIX
dns-server 2001:DB8:100::y
!
interface GigabitEthernet0/0/0
ipv6 dhcp server Customers
```

y should be substituted with 2 or 4, depending on your team assignment.

Step 3 Enter these commands on the CE router, running IOS.

```
CEy(config)#
interface GigabitEthernet0/0
ipv6 dhcp client pd SP_ASSIGNED_PREFIX
```

y should be substituted with 2 or 4, depending on your team assignment.

Step 4 Enter these commands on the CE router, running IOS.

```
CEy(config)#
interface GigabitEthernet0/1
ipv6 address SP_ASSIGNED_PREFIX ::1/64
```

y should be substituted with 2 or 4, depending on your team assignment.

Task 2: Configure DHCPv6 Lite Server

Step 1 Enter these commands on the CE router, running IOS.

```
CEy(config)#
ipv6 dhcp pool POOL
import dns-server
```

y should be substituted with 2 or 4, depending on your team assignment.

Step 2 Enter these commands on the CE router, running IOS.

```
CEy(config)#  
interface GigabitEthernet0/1  
  ipv6 nd other-config-flag  
  ipv6 dhcp server POOL
```

y should be substituted with 2 or 4, depending on your team assignment.

Step 3 Enter these commands on the CE router, running IOS.

```
CEy(config)#  
interface GigabitEthernet0/1  
  no ipv6 address SP_ASSIGNED_PREFIX ::1/64
```

y should be substituted with 2 or 4, depending on your team assignment.

Step 4 Enter these commands on the CE router, running IOS.

```
CEy(config)#  
interface GigabitEthernet0/1  
  ipv6 address autoconfig
```

y should be substituted with 2 or 4, depending on your team assignment.

Hardware Lab 9: Implement IPv6 Multicasting

Task 1: Create a New Loopback Interface and Verify Connectivity

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

```
RP/0/RSP0/CPU0:PEx(config)#  
interface Loopbackx0  
!  
commit
```

```
PEy(config)#  
interface Loopbacky0
```

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

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

```
RP/0/RSP0/CPU0:PEx(config)#  
interface Loopbackx0  
ipv6 address 2001:db8:x:x::1/128  
!  
commit
```

```
PEy(config)#  
interface Loopbacky0  
ipv6 address 2001:DB8:y:y::1/128
```

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

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

```
RP/0/RSP0/CPU0:PEx(config)#  
router ospfv3 1  
area 0  
interface Loopback x0  
!  
commit
```

```
PEy(config)#  
interface Loopbacky0  
ipv6 ospf 1 area 0
```

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

Task 2: Implement IPv6 Multicast Using Embedded RPs

Step 1 Enter these commands on the CE router, running IOS.

```
CEx(config)#  
ipv6 multicast-routing  
ipv6 pim spt-threshold infinity
```

```
CEy(config)#  
ipv6 multicast-routing  
ipv6 pim spt-threshold infinity
```

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

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

```
RP/0/RSP0/CPU0:PEx(config)#  
multicast-routing  
address-family ipv6  
interface all enable  
!  
router pim  
address-family ipv6  
spt-threshold infinity  
interface Loopbackx0  
enable  
!  
interface GigabitEthernet0/0/0/0  
enable  
!  
interface GigabitEthernet0/0/0/1  
enable  
!  
commit
```

```
PEy(config)#  
ipv6 pim spt-threshold infinity  
!  
interface Loopbacky0  
ipv6 pim  
!  
interface GigabitEthernet0/0/0  
ipv6 pim  
!  
interface GigabitEthernet0/0/1  
ipv6 pim  
!
```

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

Step 5 The multicast group addresses that can be served by the RP assigned to your pod:

Router	RP Address	Group Addresses
PE _x	2001.db8:x:x::1/128	FF75:0140:2001:db8:x:x::/96
PE _y	2001.db8:y:y::1/128	FF75:0140:2001:db8:y:y::/96

Step 6 Enter these commands on the PE router, running IOS XR and IOS XE, respectively.

```
RP/0/RSP0/CPU0:PEx(config)#
ipv6 access-list MCAST_PODx
permit ipv6 any ff75:140:2001:db8:x:x::/96
!
router pim
address-family ipv6
embedded-rp 2001:db8:x:x::1 MCAST_PODx
!
commit
```

```
PEy(config)#
ipv6 pim rp-address 2001:DB8:y:y::1
```

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

Step 7 $2^{32} = 4294967296$

Step 8 Enter these commands on the CE router, running IOS.

```
CEx(config)#
interface Loopback0
ipv6 mld join-group FF75:140:2001:DB8:x:x:0:1
```

```
CEy(config)#
interface Loopback0
ipv6 mld join-group FF75:140:2001:DB8:y:y:0:1
```

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

Hardware Lab 10: Implement Tunnels for IPv6

Task 1: Configure a Static IPv6-in-IPv4 Tunnel

Step 3 Enter these commands on the CE router, running IOS.

```
CEx(config)#  
no ip pim rp-address 1.1.1.1  
!  
interface Tunnel0  
ipv6 enable
```

```
CEy(config)#  
no ip pim rp-address 1.1.1.1  
!  
interface Tunnel0  
ipv6 enable
```

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

Step 4 Enter these commands on the CE router, running IOS.

```
CEx(config)#  
interface Tunnel0  
tunnel source Loopback0  
tunnel destination 10.y.10.1
```

```
CEy(config)#  
interface Tunnel0  
tunnel source Loopback0  
tunnel destination 10.x.10.1
```

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

Step 5 Enter these commands on the CE router, running IOS.

```
CEx(config)#  
interface Tunnel0  
tunnel mode ipv6ip
```

```
CEy(config)#  
interface Tunnel0  
tunnel mode ipv6ip
```

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

Step 6 Enter these commands on the CE router, running IOS.

```
CEx(config)#  
ipv6 route 2001:DB8:10:y:10::1/128 Tunnel0
```

```
CEy(config)#  
ipv6 route 2001:DB8:10:x:10::1/128 Tunnel0
```

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

Step 10 Enter these commands on the CE router, running IOS.

```
CEx(config)#  
no ipv6 route 2001:DB8:10:y:10::1/128 Tunnel0  
no interface Tunnel0
```

```
CEy(config)#  
no ipv6 route 2001:DB8:10:x:10::1/128 Tunnel0  
no interface Tunnel0
```

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

Task 2: Configure Dynamic 6RD Tunnels

Step 1 Enter these commands on the PE router, running IOS XE.

```
PEy(config)#  
router bgp 64500  
address-family ipv4  
network 10.y.1.1 mask 255.255.255.255
```

y should be substituted with 2 or 4, depending on your team assignment.

Step 3 The prefix is 8 bits long.

Step 4 24 bits of the IPv4 address will be used to construct 6RD networks.

Step 5 These are 6RD networks for IPv6 sites behind the CE routers:

CE Router	IPv4 Tunnel Endpoint	6RD Prefix	6RD Network
CEx	10.x.10.1	2001:db8:aa00::/40	2001:DB8:AA0x:A01::/64
CEy	10.y.10.1	2001:db8:aa00::/40	2001:DB8:AA0y:A01::/64

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

Step 6 This is the 6RD network for IPv6 sites behind the PE (Cisco IOS XE) router:

PE Router	IPv4 Tunnel Endpoint	6RD Prefix	6RD Network
PEy	10.y.1.1	2001:db8:aa00::/40	2001:DB8:AA0y:101::/64

y should be substituted with 2 or 4, depending on your team assignment.

Step 7 Enter these commands on the CE router, running IOS.

```
CEx(config)#  
interface Tunnel0  
  ipv6 enable  
  tunnel source Loopback0  
  tunnel mode ipv6ip 6rd  
  tunnel 6rd ipv4 prefix-len 8  
  tunnel 6rd prefix 2001:DB8:AA00::/40  
  tunnel 6rd br 10.y.1.1
```

```
CEy(config)#  
interface Tunnel0  
  ipv6 enable  
  tunnel source Loopback0  
  tunnel mode ipv6ip 6rd  
  tunnel 6rd ipv4 prefix-len 8  
  tunnel 6rd prefix 2001:DB8:AA00::/40  
  tunnel 6rd br 10.y.1.1
```

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

Step 9 Enter these commands on the CE router, running IOS.

```
CEx(config)#  
interface Loopback10  
  ipv6 address 2001:DB8:AA0x:A01::1/64  
  ipv6 enable
```

```
CEy(config)#  
interface Loopback10  
  ipv6 address 2001:DB8:AA0y:A01::1/64  
  ipv6 enable
```

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

Step 10 Enter these commands on the CE router, running IOS.

```
CEx(config)#  
ipv6 route 2001:DB8:AA00::/40 Tunnel0
```

```
CEy(config)#  
ipv6 route 2001:DB8:AA00::/40 Tunnel0
```

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

Step 11 Enter these commands on the PE router, running IOS XE.

```
PEy(config)#  
interface Tunnel10  
ipv6 enable  
tunnel source Loopback0  
tunnel mode ipv6ip 6rd  
tunnel 6rd ipv4 prefix-len 8  
tunnel 6rd prefix 2001:DB8:AA00::/40
```

y should be substituted with 2 or 4, depending on your team assignment.

Step 12 Enter these commands on the PE router, running IOS XE.

```
PEy(config)#  
ipv6 route 2001:DB8:AA00::/40 Tunnel10
```

y should be substituted with 2 or 4, depending on your team assignment.

Step 14 Enter these commands on the PE router, running IOS XE.

```
PEy(config)#  
interface Loopback10  
ipv6 address 2001:DB8:101:y::1/64  
ipv6 enable
```

y should be substituted with 2 or 4, depending on your team assignment.

Step 15 Enter these commands on the CE router, running IOS.

```
CEx(config)#  
ipv6 route ::/0 Tunnel0 2001:DB8:AA0y:101::
```

```
CEy(config)#  
ipv6 route ::/0 Tunnel0 2001:DB8:AA0y:101::
```

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