360RSW05-WKB-CA09

Cisco 360 CCIE R&S Exercise Workbook Introduction

The Cisco 360 CCIE® R&S Exercise Workbook contains 20 challenging scenarios at the CCIE level that can be used for rigorous self-paced practice.

Each lab provides an extensive answer key, Mentor Guide support, and verification tables and is designed to maximize learning by providing practical experience. Also, self-paced learning resources such as the Cisco 360 CCIE R&S Reference Library and Cisco 360 CCIE R&S lessons supplement the Exercise Workbook scenarios.

Cisco 360 CCIE R&S Exercise Workbook Lab 9 Configuration Section Answer Key

Table of Contents

Answer Key Structure

Section One

The answer key PDF document is downloadable from the web portal.

Section Two

To obtain a comprehensive view of the configuration for a specific section, access the Mentor Guide engine in the web portal.

Exercise Workbook Lab 9 Configuration Section Answer Key

Note	Regardless of any configuration you perform in this lab, it is very important that you conform to the general guidelines that are provided in the "Restrictions and Goals" section. If you do not conform to the guidelines, you could have a significant deduction of points in your final score.

Grading and Duration

- Configuration lab duration: 6 hours
- Configuration lab maximum score: 76 points

Note	You can assess your progress on the self-paced labs in this workbook by adding up the points that are assigned to sections and tasks. Consider taking the full Assessment Labs to assess your readiness level.

Difficulty Level

- Difficulty: Basic to Intermediate

Restrictions and Goals

Note	Read this section carefully.

- To receive any credit for a subsection, you must fully complete the subsection as per requirements. You will *not* receive partial credit for partially completed subsections.
- IPv4 subnets displayed in the scenario diagram belong to network 148.49.0.0/16.
- *Points will be deducted from multiple sections for failing to assign correct IPv4 addresses.*
- Do not use any static routes.
- Advertise loopback interfaces with their original masks.
- Network 0.0.0.0/0 should not appear in any routing table (**show ip route**) except on R8.
- Do not use the **ip default-gateway** or **ip default-network** commands.
- Do not introduce any new IP addresses.
- All IP addresses involved in this scenario must be reachable, unless explicitly specified otherwise.

- Unless explicitly specified otherwise, addresses and networks that are advertised in the "Border Gateway Protocol" (BGP) section need to be reachable by all BGP routers but do not have to be reachable by interior gateway protocol (IGP)-only routers.
- Do not create new interfaces to fulfill IGP requirements, and do not create any summaries, unless the summary is required to meet explicitly stated scenario requirements.
- Do not introduce any new IPv4 or IPv6 addresses unless the instructions explicitly specify otherwise.
- Use only conventional routing algorithms, unless specified otherwise.
- Do not modify the hostname, console, or vty configuration unless you are specifically asked to do so.
- Do not modify the initial interface or IP address numbering.

Explanation of Each of the Restrictions and Goals

IPv4 subnets that are displayed in the scenario "IPv4 IGP" diagram belong to network 148.49.0.0/16.
All IP addresses in this lab belong to the 148.49.0.0/16 address space, except for prefixes that are explicitly specified to be part of a different IP space.

Do not use any static routes.

Static routes can be used to solve a range of reachability problems. However, you cannot use them in this lab. You must rely on skillful configuration of all your unicast routing protocols.

Advertise loopback interfaces with their original masks.

The original mask is the mask configured on the loopback interface. OSPF treats loopback interfaces as host routes by default and advertises them as /32 prefixes. The requirement to advertise loopback interfaces with their original masks precludes using the default OSPF network type for the loopback interface. You need to provide a solution, such as changing the OSPF network type or summarizations.

Network 0.0.0.0/0 should not appear in any routing table (show ip route**).**

A 0.0.0.0/0 entry can be used to solve a range of reachability problems. In particular, a 0.0.0.0/0 entry can be used to set up the gateway of last resort. In this exercise, you cannot use any 0.0.0.0/0 entries. Route summarization is an alternative to using the 0.0.0.0/0 route to solve the reachability problem. R8 is excluded from this requirement.

Do not use the ip default-network **command.**

This command can be used to solve reachability issues by setting the gateway of last resort. This command generates 0.0.0.0/0 in the Routing Information Protocol (RIP) environment. You cannot use it in this scenario.

All the IP addresses that are involved in this scenario must be reachable.

This goal is a key goal to observe. It requires that all your IGPs and your routing policy tasks be configured properly. The key elements of your routing policy include route redistribution and the controlling of routing updates using the **distribute-list**, **route-map**, and **distance** commands. A key point to remember about this lab is that the term "redistribution" is not explicitly used. However, you must perform redistribution to ensure that all IP addresses are reachable without the use of static routes or 0.0.0.0/0 routes.

Addresses and networks that are advertised in the BGP section need to be reachable by all BGP routers but do not have to be reachable by IGP-only routers.

This statement relaxes the requirement that all IP addresses must be reachable. The BGP prefixes need only be reachable among the routers specified in the BGP section. They can be used in other unicast tables. However, BGP routers need to have the prefixes in the routing tables and need to be able to forward traffic to the addresses known via BGP.

Use conventional routing algorithms.

This restriction prevents you from solving any problems by configuring policy routing. At the heart of this restriction is the interpretation of "conventional routing algorithms." Although this phrase can be interpreted in different ways, this interpretation is applied in this workbook:

> Conventional routing algorithms are routing algorithms that apply destination-based prefix lookups in a routing table. Conventional routing algorithms do not use any type of information other than the destination address to make a packet forwarding decision.

Because of this restrictive interpretation, no form of policy routing can be applied. Whenever you see this restriction, you will need to use dynamic routing protocols to fulfill all packet-forwarding requirements.

1. Switch Configuration

General Tasks:

As with any switch configuration, you must address the following basic configuration requirements: setting the VLAN Trunking Protocol (VTP) mode, configuring trunk ports, and statically assigning ports to VLANs. For a good reference on mastering basic Cisco Catalyst 3560 Switch configuration tasks, access the full set of Catalyst video-on-demand (VoD) sessions within the "Link Layer" lesson in the Cisco 360 learning portal. These self-paced sessions provide more than 7 hours of instruction on a range of basic Catalyst switch configuration tasks.

Note that not all Cisco Catalyst 3560 Switch configuration features are supported on the virtual Cisco IOS Software on UNIX.

Configure the VLANs and the VLAN names according to the scenario specifications and assign the ports of the switches to these VLANs. Make sure that the VLAN names are spelled correctly and match the letter case.

Use the "VLAN Allocation" and "Switch-to-Switch Connections" tables to analyze the VLAN propagation in this lab.

See the following diagram for the VLAN layout.

VLAN Propagation

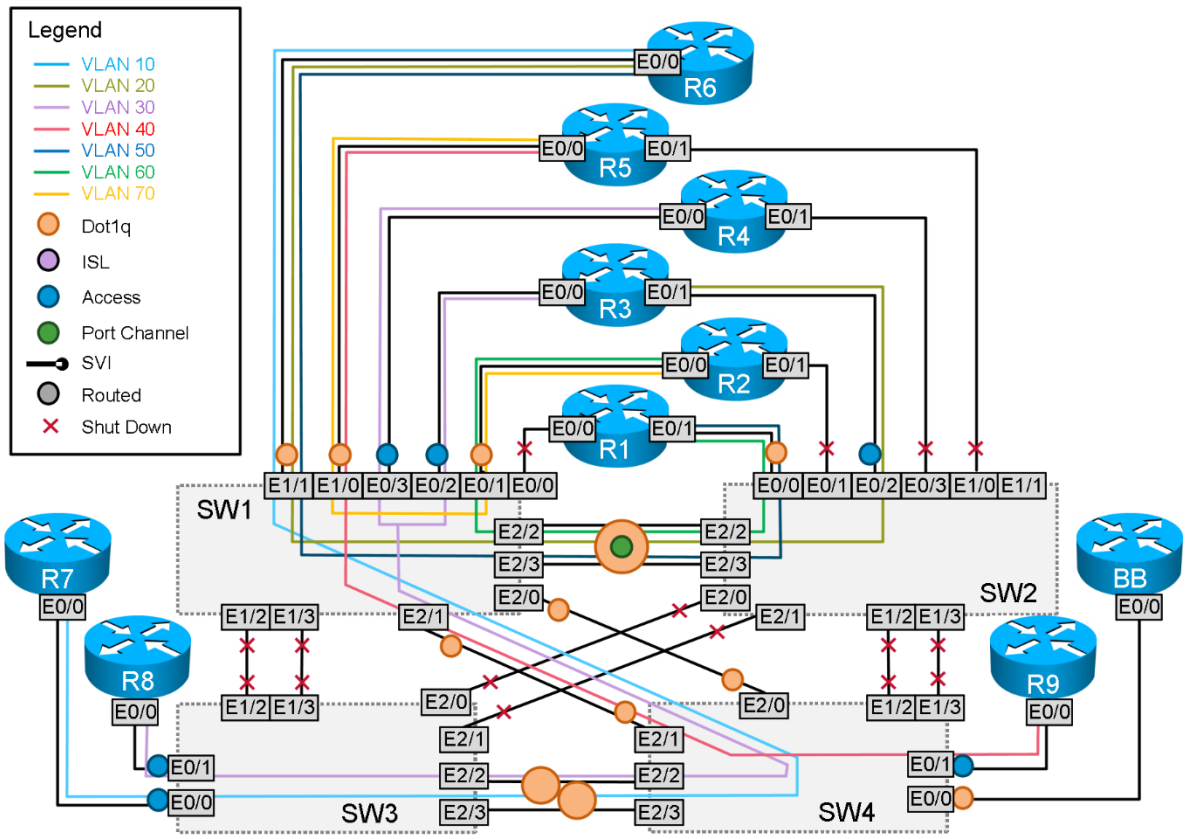

Carefully review the entire scenario. Closely examine the supplied diagram and any associated tables. Determine how you need to configure VTP, how to configure ports that are assigned as trunks, and how to configure ports that are assigned as static VLAN ports. Use the **switchport mode access** command to statically assign ports to a VLAN.

Issue: No VLANs should be configured in the configuration mode on SW2, SW3, or SW4. All switches are in the TEST domain.

Solution:

The VLAN distribution diagram indicates that SW2, SW3, and SW4 should have VLANs configured. How can you configure VLANs on a switch without explicitly configuring the VLANs in the switch's configuration mode? If you configure switches SW2, SW3, and SW4 as VTP clients, SW1 must be configured as a VTP server. With SW1 configured as a VTP server, you can configure the VLANs on SW1 and advertise them to the clients (in this scenario, the VTP clients are SW2, SW3, and SW4). To fulfill all the configuration requirements, change the name of the VTP domain using the **vtp domain TEST** command.

SW1:

```
vtp domain TEST
```

SW2:

```
vtp mode client
```

SW3:

```
vtp mode client
```

SW4:

```
vtp mode client
```

Issue: Set VTP version 2 on switch SW2.

Solution:

If the switch is in VTP client mode, you cannot change the VTP version:

```
SW2(config)#vtp version 2
Cannot modify version in VTP client mode unless the system is in VTP version 3
SW2(config)#
```

You must go to the VTP server mode and change the version there.

When switches are in the VTP server-client mode, the specific VTP and VLAN information pertaining to this configuration is stored in the **vlan.dat** file, and not in the startup configuration file **config.text**.

```
SW1#show vtp status
VTP Version                     : 3 (capable)
Configuration Revision          : 8
Maximum VLANs supported locally : 1005
Number of existing VLANs        : 12
VTP Operating Mode              : Server
VTP Domain Name                 : TEST
VTP Pruning Mode                : Disabled (Operationally Disabled)
VTP V2 Mode                     : Enabled
VTP Traps Generation            : Disabled
MD5 digest                      : 0x5A 0x5F 0x6E 0x41 0x4B 0x58 0xEB 0x17
Configuration last modified by 0.0.0.0 at 8-9-13 19:49:00
```



```
Local updater ID is 0.0.0.0 (no valid interface found)
VTP version running             : 2
SW1#
SW2#show vtp status
VTP Version                     : 3 (capable)
Configuration Revision          : 8
Maximum VLANs supported locally : 1005
Number of existing VLANs        : 12
VTP Operating Mode              : Client
VTP Domain Name                 : TEST
VTP Pruning Mode                : Disabled (Operationally Disabled)
VTP V2 Mode                     : Enabled
VTP Traps Generation            : Disabled
MD5 digest                      : 0x5A 0x5F 0x6E 0x41 0x4B 0x58 0xEB 0x17
Configuration last modified by 0.0.0.0 at 8-9-13 19:49:00
VTP version running             : 2
SW2#
SW3#show vtp status
VTP Version                     : 3 (capable)
Configuration Revision          : 8
Maximum VLANs supported locally : 1005
Number of existing VLANs        : 12
VTP Operating Mode              : Client
VTP Domain Name                 : TEST
VTP Pruning Mode                : Disabled (Operationally Disabled)
VTP V2 Mode                     : Enabled
VTP Traps Generation            : Disabled
MD5 digest                      : 0x5A 0x5F 0x6E 0x41 0x4B 0x58 0xEB 0x17
Configuration last modified by 0.0.0.0 at 8-9-13 19:49:00
VTP version running             : 2
SW3#
SW4#show vtp status
VTP Version                     : 3 (capable)
Configuration Revision          : 8
Maximum VLANs supported locally : 1005
Number of existing VLANs        : 12
VTP Operating Mode              : Client
VTP Domain Name                 : TEST
VTP Pruning Mode                : Disabled (Operationally Disabled)
VTP V2 Mode                     : Enabled
VTP Traps Generation            : Disabled
MD5 digest                      : 0x5A 0x5F 0x6E 0x41 0x4B 0x58 0xEB 0x17
Configuration last modified by 0.0.0.0 at 8-9-13 19:49:00
VTP version running             : 2
SW4#
```

Issue: Ensure that SW1 is the root bridge for VLAN 30.
Ensure that SW2 is the root bridge for VLAN 20 and VLAN 40.
Ensure that SW4 is the root bridge for VLAN 10.

Solution:

The root bridge election happens on a per-VLAN basis. You can change the default bridge priority to a lower number for a particular VLAN to ensure that the desired switch is elected as the root bridge:

```
SW1(config)#spanning-tree vlan 30 root primary
SW2(config)#spanning-tree vlan 20,40 root primary
SW4(config)#spanning-tree vlan 10 root primary

SW1#sh spanning-tree vlan 30 root

                                        Root Hello Max Fwd
Vlan                Root ID             Cost  Time Age Dly  Root Port
---------------- ------------------- ------ ----- --- ---  ----------------
```

```
VLAN0030         24606 aabb.cc00.0700         0    2   20  15
SW1#

SW2#show spanning-tree vlan 20,40 root

                                        Root Hello Max Fwd
Vlan                 Root ID            Cost  Time Age Dly  Root Port
---------------- -------------------- ------ ----- --- ---  ----------------
VLAN0020         24596 aabb.cc00.0800         0    2   20  15
VLAN0040         24616 aabb.cc00.0800         0    2   20  15
SW2#

SW4#show spanning-tree vlan 10 root

                                        Root Hello Max Fwd
Vlan                 Root ID            Cost  Time Age Dly  Root Port
---------------- -------------------- ------ ----- --- ---  ----------------
VLAN0010         24586 aabb.cc00.0a00         0    2   20  15
SW4#
```

Notice that the root cost is always 0 on the root bridge.

```
SW1#sh spanning-tree root

                                        Root Hello Max Fwd
Vlan                 Root ID            Cost  Time Age Dly  Root Port
---------------- -------------------- ------ ----- --- ---  ----------------
VLAN0010         24586 aabb.cc00.0a00       100    2   20  15  Et2/0
VLAN0020         24596 aabb.cc00.0800        56    2   20  15  Po1
VLAN0030         24606 aabb.cc00.0700         0    2   20  15
VLAN0040         24616 aabb.cc00.0800        56    2   20  15  Po1
VLAN0050         32818 aabb.cc00.0700         0    2   20  15
VLAN0060         32828 aabb.cc00.0700         0    2   20  15
VLAN0070         32838 aabb.cc00.0700         0    2   20  15
SW1#
```

On SW1, the root cost for VLAN 20 and VLAN 40 is 56, because SW2 is the root bridge for these VLANs and 56 is the cost of the EtherChannel logical interface.

Issue: Configure forwarding on VLAN 10 to prefer the link between ports 2/1 of SW1 and SW4. If the link between ports 2/1 of SW1 and SW4 fails, forward traffic through the link between ports 2/0 of SW1 and SW4.

Solution:

There are two links between SW1 and SW4: between ports Et2/1 and between ports Et2/0. Both links are dot1q trunks and VLAN 10 is allowed on both trunks. Spanning tree will block one of the trunk links:

```
SW1#show spanning-tree vlan 10

VLAN0010
  Spanning tree enabled protocol ieee
  Root ID     Priority    24586
              Address     aabb.cc00.0a00
              Cost        100
              Port        65 (Ethernet2/0)
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID   Priority    32778  (priority 32768 sys-id-ext 10)
              Address     aabb.cc00.0700
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
              Aging Time 300

Interface           Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- --------------------------------
```



```
Et1/1               Desg FWD 100       128.34   Shr
Et2/0               Root FWD 100       128.65   Shr
Et2/1               Altn BLK 100       128.66   Shr

SW1#
```

With the root bridge on SW4, the blocked port is Et2/1 by default, because the root path costs are equal and the port ID sent by SW4 Et2/1 is greater than the port ID sent by SW4 Et2/0.

To have SW1 choose Et2/1 as its root port, you could make its root path cost or received port ID lower than SW1 Et2/0. The path cost manipulation is chosen in this answer key:

```
SW1(config)#int e2/1
SW1(config-if)#spanning-tree vlan 10 cost 20


SW1#show spanning-tree vlan 10

VLAN0010
  Spanning tree enabled protocol ieee
  Root ID     Priority    24586
              Address     aabb.cc00.0a00
              Cost        20
              Port        66 (Ethernet2/1)
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID   Priority    32778  (priority 32768 sys-id-ext 10)
              Address     aabb.cc00.0700
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
              Aging Time 15

Interface           Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- --------------------------------
Et1/1               Desg FWD 100       128.34   Shr
Et2/0               Altn BLK 100       128.65   Shr
Et2/1               Root FWD 20        128.66   Shr

SW1#
```

This fulfills the requirements of this task.

Note	To obtain a comprehensive view of the configuration tasks in this section, access the Mentor Guide engine. With the Mentor Guide engine, you can enter more than 1000 Cisco IOS Software commands, as well as a collection of proprietary commands such as **show all**.

2. IPv4 OSPF

Note	All OSPF routers must be configured with only one OSPF Process ID (PID). Use your IGP diagram to help guide configuration.

Issue: Configure OSPF Area 0 between routers R1, R3, and R4 on subnet 148.49.134.0/24. Do not use the default OSPF network type on the OSPF Area 0 interfaces. Form the Area 0 adjacency without using the **neighbor** command.

Solution:

The interfaces on subnet 148.49.134.0/24 are Ethernet; therefore, the OSPF network type defaults to broadcast. However, the default network type of broadcast cannot be used according to the lab requirements.

These are the options for OSPF network type configuration:

- broadcast
- nonbroadcast
- point-to-multipoint
- point-to-multipoint nonbroadcast
- point-to-point

The point-to-point option can be eliminated because it cannot be used with more than two shared interfaces.

The lab requirements state that the OSPF adjacency must be formed without the **neighbor** statement. This means that you cannot use the nonbroadcast or point-to-multipoint nonbroadcast network types, because these OSPF network types require the **neighbor** configuration.

Therefore, only the point-to-multipoint network type can be used in this lab for the interfaces connected to the subnet 148.49.134.0/24.

Here is an example of the OSPF configuration on R1:

```
R1#show ip ospf interface brief
Interface    PID   Area            IP Address/Mask    Cost  State Nbrs F/C
Et0/2        1     0               148.49.134.1/24    10    P2MP  2/2
R1#
```

Issue: Configure OSPF Area 30 between routers R3, R4, and R8 on VLAN 30 subnet 148.49.34.0/24.
Use the default OSPF network type to form the Area 30 adjacency.
Router R3 must be elected as the DR.
No BDR must be elected on VLAN 30.

Solution:

R3, R4, and R8 are on an Ethernet broadcast VLAN 30 segment. The default OSPF network type for the Ethernet interfaces is broadcast. To ensure that there is no BDR on this segment, configure the DR on R3 and DROTHER (any router that is neither DR nor BDR) on R8 and R4 by setting the OSPF priority to 0 on the interfaces of R4 and R8 that are connected to VLAN 30. Since the OSPF network type is broadcast, hello packets will be sent to the multicast destination. Therefore, the **neighbor** statements are not required.

```
R3#show ip ospf neighbor e0/0 detail
 Neighbor 148.49.108.1, interface address 148.49.34.8
    In the area 30 via interface Ethernet0/0
    Neighbor priority is 0, State is FULL, 6 state changes
    DR is 148.49.34.3 BDR is 0.0.0.0
    Options is 0x18 in Hello (N/P-bit, L-bit)
    Options is 0x58 in DBD (N/P-bit, L-bit, O-bit)
    LLS Options is 0x1 (LR)
    Dead timer due in 00:00:33
    Neighbor is up for 3d14h
    Index 2/4, retransmission queue length 0, number of retransmission 0
    First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)
    Last retransmission scan length is 0, maximum is 0
    Last retransmission scan time is 0 msec, maximum is 0 msec
 Neighbor 148.49.255.4, interface address 148.49.34.4
    In the area 30 via interface Ethernet0/0
    Neighbor priority is 0, State is FULL, 6 state changes
    DR is 148.49.34.3 BDR is 0.0.0.0
    Options is 0x18 in Hello (N/P-bit, L-bit)
    Options is 0x58 in DBD (N/P-bit, L-bit, O-bit)
```



```
    LLS Options is 0x1 (LR)
    Dead timer due in 00:00:32
    Neighbor is up for 3d18h
    Index 1/2, retransmission queue length 0, number of retransmission 1
    First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)
    Last retransmission scan length is 1, maximum is 1
    Last retransmission scan time is 0 msec, maximum is 0 msec
R3#
```

Issue: R3 should use 148.49.255.3 as its OSPF router ID.

Solution:

By default, the router selects its highest loopback address for its OSPF router ID.

You can manually override this behavior with this input on R3:

```
router ospf 1
router-id 148.49.255.3
```

R3:

```
R3#show ip ospf | inc ID
 Routing Process "ospf 1" with ID 148.49.255.3
R3#
```

That sets the stage for the task that requires defining the router that will be the Type-7-to-Type-5 Translator for the not-so-stubby-area (NSSA).

Issue: No external routing information is allowed into OSPF Area 30 from the backbone area. However, advertise the Loopback108 network on R8 as an OSPF external prefix. The network of the Loopback108 interface must be listed in the routing tables of both R3 and R4 as an "N1" OSPF prefix.

Solution:

A recommended starting point for fulfilling the configuration requirements above is to inject the Loopback108 interface residing on R8 into OSPF as an external prefix. This can only be performed by configuring **redistribute connected** on R8.

Once this operation has been performed, it excludes OSPF Area 30 from being configured as stub or totally stubby area. However, it does not exclude OSPF Area 30 from being configured as an NSSA. The lab configuration requirement that "no external routing information is allowed into OSPF Area 30 from the backbone area" means that you should configure OSPF Area 30 as an NSSA to block the external link-state advertisements (LSAs) from being propagated from OSPF Area 0 into Area 30, as well as to allow external information from R8 (the Loopback108 **redistributed connected** network).

R3:

```
router ospf 1
 area 30 nssa
 redistribute connected subnets route-map Loopback108
 network 148.49.34.8 0.0.0.0 area 30
!
!
route-map Loopback108 permit 10
 match interface Loopback108
 set metric-type type-1
!
R8#show ip ospf database nssa-external self-originate

            OSPF Router with ID (148.49.108.1) (Process ID 1)
```

```
                Type-7 AS External Link States (Area 30)

  LS age: 360
  Options: (No TOS-capability, Type 7/5 translation, DC, Upward)
  LS Type: AS External Link
  Link State ID: 148.49.108.0 (External Network Number )
  Advertising Router: 148.49.108.1
  LS Seq Number: 800000A3
  Checksum: 0x25F9
  Length: 36
  Network Mask: /24
        Metric Type: 1 (Comparable directly to link state metric)
        MTID: 0
        Metric: 20
        Forward Address: 148.49.34.8
        External Route Tag: 0

R8#
```

Issue: R1 must use the R4 router as the next hop to reach the network 148.49.108.0/24. Do not use any filtering technique to accomplish this task.

Solution:

Closely examine the following diagram:

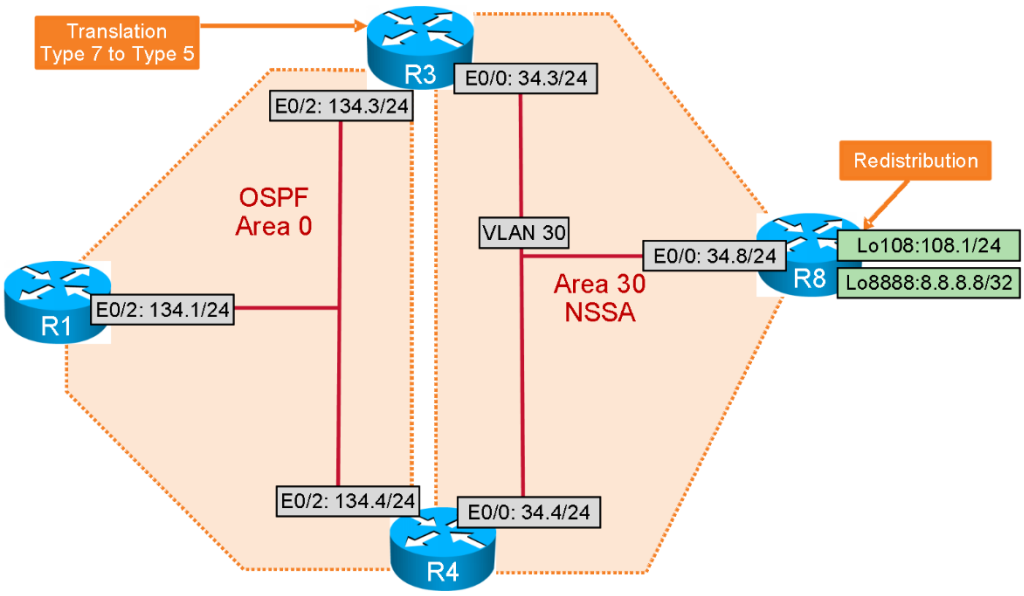

OSPF Area 30 has two Area Border Routers (ABRs) in this topology (R3 and R4); however, only one ABR will perform the translation of LSA Type 7 into LSA Type 5. The translating ABR will be selected as the OSPF ABR router with the highest router ID. Because router ID 148.49.255.3 has been configured on R3, R3 will have a higher loopback IP address than R4; therefore, R3 will perform the translation. In other words, R3 will originate the LSA Type 5 for the 148.49.108.0/24 network into the backbone Area 0. Therefore, R1 will use R3 as the next hop for the Loopback108 interface originating from R8. Using R4 as the next hop would contradict the scenario requirement.

The solution is to choose the IP addresses for the router ID on R4 so that the IP address of R3 is lower than the IP address on R4. The IP address used for the OSPF router ID does not have to be configured on an interface of the router. Since R3 must have 148.49.255.3 as an OSPF router ID, 148.49.255.4 can be configured as the OSPF router ID on R4.

```
R4#show ip ospf | inc ID
 Routing Process "ospf 1" with ID 148.49.255.4
R4#
```

The following diagram displays the Translation and Type 7, Type 5 LSA propagation:

LSA Propagation

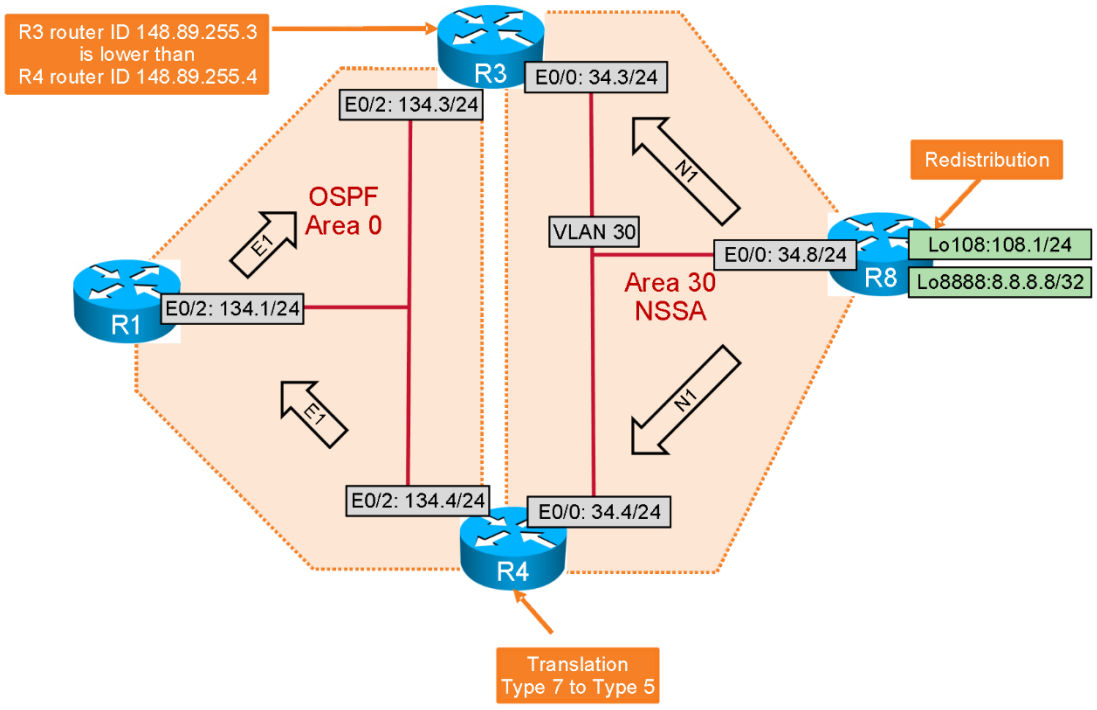

Examine the OSPF configuration and analyze the results.

Verify the OSPF adjacency table on R1:

```
R1#show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
148.49.255.4      0   FULL/  -        00:01:45    148.49.134.4    Ethernet0/2
148.49.255.3      0   FULL/  -        00:01:39    148.49.134.3    Ethernet0/2
R1#
```

Note that in Area 0, the OSPF network type is point-to-multipoint, so there is no DR.

Verify the OSPF adjacency tables on R3 and R4:

```
R3#show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
148.49.255.4      0   FULL/  -        00:01:52    148.49.134.4    Ethernet0/2
148.49.101.1      0   FULL/  -        00:01:48    148.49.134.1    Ethernet0/2
148.49.108.1      0   FULL/DROTHER    00:00:30    148.49.34.8     Ethernet0/0
148.49.255.4      0   FULL/DROTHER    00:00:32    148.49.34.4     Ethernet0/0
R3#

R4#show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
148.49.101.1      0   FULL/  -        00:01:43    148.49.134.1    Ethernet0/2
148.49.255.3      0   FULL/  -        00:01:42    148.49.134.3    Ethernet0/2
148.49.108.1      0   2WAY/DROTHER    00:00:35    148.49.34.8     Ethernet0/0
148.49.255.3      1   FULL/DR         00:00:31    148.49.34.3     Ethernet0/0
R4#
```

Verify the OSPF adjacency table on R8:

```
R8#show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
148.49.255.3      1   FULL/DR         00:00:35    148.49.34.3     Ethernet0/0
148.49.255.4      0   2WAY/DROTHER    00:00:39    148.49.34.4     Ethernet0/0
R8#
```

On VLAN 30, the OSPF network type is broadcast and R3 is the DR.

Verify the Loopback108 network in the routing tables on R1, R3, and R4.

R1:

```
R1#show ip route 148.49.108.0
Routing entry for 148.49.108.0/24
  Known via "ospf 1", distance 110, metric 30, type extern 1
  Redistributing via eigrp 200, rip
  Advertised by eigrp 200 metric 1000 100 255 2 1500 route-map OSPFtoEIGRP
                rip route-map toRIP
  Last update from 148.49.134.4 on Ethernet0/2, 3d14h ago
  Routing Descriptor Blocks:
  * 148.49.134.4, from 148.49.255.4, 3d14h ago, via Ethernet0/2
      Route metric is 30, traffic share count is 1
R1#
R1#show ip route ospf | inc 108
O E1     148.49.108.0/24 [110/40] via 148.49.134.4, 00:16:57, Ethernet0/2
R1#
```

Note that the route 148.49.108.0/24 is learned as the OSPF external type 1 (E1) prefix on R1. R1 prefers 148.49.108.0/24 via R4, because R4 performs the translation of the OSPF NSSA N1 prefixes to the OSPF external E1 prefixes.

R3:


```
R3#show ip route 148.49.108.0
Routing entry for 148.49.108.0/24
  Known via "ospf 1", distance 110, metric 30, type NSSA extern 1
  Redistributing via eigrp 200
  Advertised by eigrp 200 metric 1000 100 255 2 1500 route-map OSPFtoEIGRP
  Last update from 148.49.34.8 on Ethernet0/0, 3d14h ago
  Routing Descriptor Blocks:
  * 148.49.34.8, from 148.49.108.1, 3d14h ago, via Ethernet0/0
      Route metric is 30, traffic share count is 1
R3#
R3#show ip route ospf | inc 108
O N1      148.49.108.0/24 [110/30] via 148.49.34.8, 00:19:15, Ethernet0/0
R3#
```

Note that the route 148.49.108.0/24 is learned as the OSPF NSSA E1 prefix on R3. R3 prefers 148.49.108.0/24 via the originating router R8.

R4:

```
R4#show ip route 148.49.108.0
Routing entry for 148.49.108.0/24
  Known via "ospf 1", distance 110, metric 30, type NSSA extern 1
  Last update from 148.49.34.8 on Ethernet0/0, 3d14h ago
  Routing Descriptor Blocks:
  * 148.49.34.8, from 148.49.108.1, 3d14h ago, via Ethernet0/0
      Route metric is 30, traffic share count is 1
R4#
R4#show ip route ospf | inc 108
O N1      148.49.108.0/24 [110/30] via 148.49.34.8, 00:20:08, Ethernet0/0
R4#
R4#show ip ospf database external self-originate

            OSPF Router with ID (148.49.255.4) (Process ID 1)

                Type-5 AS External Link States

  LS age: 1447
  Options: (No TOS-capability, DC, Upward)
  LS Type: AS External Link
  Link State ID: 148.49.108.0 (External Network Number )
  Advertising Router: 148.49.255.4
  LS Seq Number: 80000001
  Checksum: 0xE351
  Length: 36
  Network Mask: /24
        Metric Type: 1 (Comparable directly to link state metric)
        MTID: 0
        Metric: 20
        Forward Address: 148.49.34.8
        External Route Tag: 0

R4#
```

Note that the route 148.49.108.0/24 is learned as the OSPF NSSA E1 prefix on R4. R4 prefers 148.49.108.0/24 via the originating router R8. R4 translates and originates the external LSA for the prefix 148.49.108.0/24.

Issue: The Loopback101 network should be advertised as an external network of type 2 with a metric of 40.

Solution:

Create the external LSA by configuring the **redistribute connected** command, referencing Loopback101 with a route map to match this specific network, and change the metric from the default 20 to the specified metric value of 40.

R1:

```
router ospf 1
redistribute connected subnets route-map LOOP101
!
route-map LOOP101 permit 10
 match interface Loopback101
 set metric 40
```

```
R1#show ip ospf database external 148.49.101.0

            OSPF Router with ID (148.49.101.1) (Process ID 1)

                Type-5 AS External Link States

  LS age: 416
  Options: (No TOS-capability, DC, Upward)
  LS Type: AS External Link
  Link State ID: 148.49.101.0 (External Network Number )
  Advertising Router: 148.49.101.1
  LS Seq Number: 80000002
  Checksum: 0xEA49
  Length: 36
  Network Mask: /24
        Metric Type: 2 (Larger than any link state path)
        MTID: 0
        Metric: 40
        Forward Address: 0.0.0.0
        External Route Tag: 0

R1#

R4#sh ip route ospf | inc 101
O E2     148.49.101.0/24 [110/40] via 148.49.134.1, 00:38:16, Ethernet0/2
R4#
```

Note that the Loopback101 interface network is advertised from R1 as the OSPF E2 prefix with a metric of 40.

Note	To obtain a comprehensive view of the configuration tasks in this section, access the Mentor Guide engine. With the Mentor Guide engine, you can enter more than 1000 Cisco IOS Software commands, as well as a collection of proprietary commands such as **show all**.

3. IPv6 OSPF

Issue: Configure OSPF IPv6 Area 0 on the link FEC0::160:0/125 between routers R1 and R6.

Solution:

The IPv6 OSPF configuration procedure is slightly different from the OSPF IPv4 configuration procedure. Here are the steps:

1. Assign the IPv6 addresses on the logical Ethernet interfaces connected to VLAN 50 of routers R1 and R6. Ensure that you can ping IPv6 between routers R1 and R6.

```
R1# show run int E0/1.50 | inc ipv6
 ipv6 address FEC0::160:1/125

R6# show run int E0/0.50 | inc ipv6
 ipv6 address FEC0::160:6/125

R1#ping ipv6 FEC0::160:6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to FEC0::160:6, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

2. Enable IPv6 unicast routing on the router with the **ipv6 unicast-routing** global configuration command. Configure the instance of the OSPF IPv6 process with a process ID (PID), for example, "1."

```
R1#sh run | inc ipv6 router ospf
ipv6 router ospf 1

R6#sh run | inc ipv6 router ospf
ipv6 router ospf 1
```

3. Assign the VLAN 50 interfaces to OSPF IPv6 Area 0 by accessing the interface configuration mode and specifying the IPv6 OSPF area configuration there:

R1:

```
interface E0/1.50
ipv6 address FEC0::160:1/125
 ipv6 ospf 1 area 0
```

R6:

```
interface E0/0.50
ipv6 address FEC0::160:6/125
 ipv6 ospf 1 area 0
```

Verify the OSPF IPv6 adjacency on R1:

```
R1#show ipv6 ospf neighbor detail

            OSPFv3 Router with ID (148.49.101.1) (Process ID 1)

 Neighbor 148.49.106.1
    In the area 0 via interface Ethernet0/1.50
    Neighbor: interface-id 18, link-local address FE80::A8BB:CCFF:FE00:600
    Neighbor priority is 1, State is FULL, 6 state changes
    DR is 148.49.106.1 BDR is 148.49.101.1
    Options is 0x000013 in Hello (V6-Bit, E-Bit, R-bit)
    Options is 0x000013 in DBD (V6-Bit, E-Bit, R-bit)
    Dead timer due in 00:00:36
    Neighbor is up for 00:52:26
    Index 1/1/1, retransmission queue length 0, number of retransmission 2
    First 0x0(0)/0x0(0)/0x0(0) Next 0x0(0)/0x0(0)/0x0(0)
    Last retransmission scan length is 1, maximum is 1
    Last retransmission scan time is 0 msec, maximum is 0 msec
R1#
R1#show ipv6 ospf | inc ID
```

```
 Routing Process "ospfv3 1" with ID 148.49.101.1
R1#
```

By default, OSPF IPv6 uses the highest IPv4 loopback address for the router ID.

Tip	OSPF IPv6 interface and area association is done on the interface level.

Issue: Place the loopback networks FEC0::101:0/125 and FEC0::106:0/125 in the OSPF IPv6 Areas 101 and 106, respectively, on routers R1 and R6.

Solution:

1. Advertise the loopback networks in Areas 101 and 106 on routers R1 and R6:

```
interface Loopback101
 ip address 148.49.101.1 255.255.255.0
 ipv6 address FEC0::101:1/125
 ipv6 ospf 1 area 101
!
interface Loopback106
 ip address 135.15.106.1 255.255.255.0
 ipv6 address FEC0::106:1/125
 ipv6 ospf 1 area 106
```

2. Verify the OSPF IPv6 networks in the IPv6 routing table:

```
R1#show ipv6 route ospf
IPv6 Routing Table - 15 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP
       U - Per-user Static route
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
       D - EIGRP, EX - EIGRP external
OI  FEC0::106:1/128 [110/1]
     via FE80::216:9DFF:FE43:F0D1, Ethernet0/1.50
```

Notice that the IPv6 loopback is advertised as a host /128 entry. This is similar to IPv4. According to the "Restrictions and Goals" section, you must change the default OSPF IPv6 advertising behavior that is applied to loopback interfaces. You also have to configure the **ipv6 ospf network point-to-point** command under the loopback interface:

```
R1#show ipv6 rout ospf
IPv6 Routing Table - 15 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP
       U - Per-user Static route
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
       D - EIGRP, EX - EIGRP external
OI  FEC0::106:0/125 [110/2]
     via FE80::216:9DFF:FE43:F0D1, Ethernet0/1.50
```

Note	To obtain a comprehensive view of the configuration tasks in this section, access the Mentor Guide engine. With the Mentor Guide engine, you can enter more than 1000 Cisco IOS Software commands, as well as a collection of proprietary commands such as **show all**.

4. IPv4 RIP

Issue: Configure a RIP routing information exchange between routers R1, R2, and R5. Updates should be sent only on VLAN 60, on the 148.49.21.0/24 subnet, and on VLAN 70.

Solution:

You must use the **passive-interface** command to control RIP update propagation.

The recommended course is to use the **passive-interface default** command to put all the interfaces that match the major networks specified in the RIP process to a passive mode (listen but do not talk). Use **no passive-interface** *interface* for the interfaces that you want to participate in the RIP updates exchange.

Issue: Ensure that all interfaces that are exchanging RIP updates send and receive RIP version 2 packets only.
Do not use the interface configuration mode to accomplish this task.

Solution:

There are two options to control the version of RIP updates sent and received on a router:

- The first option is the interface configuration mode option, where you can control the version with the **ip rip send version** and **ip rip receive version** commands. However, this option is not allowed because of the constraints of this configuration task.
- The permitted option is to specify the version for the entire RIP process under the router RIP configuration.

R1:

```
router rip
 version 2
 passive-interface default
 no passive-interface Ethernet0/1.60
 no passive-interface Serial1/0
 network 148.49.0.0
```

R2:

```
router rip
 version 2
 passive-interface default
 no passive-interface Serial1/0
 no passive-interface Ethernet0/0.60
 no passive-interface Ethernet0/0.70
 network 148.49.0.0
```

R5:

```
router rip
 version 2
 passive-interface default
 no passive-interface Ethernet0/0.70
 network 148.49.0.0
```

Issue: R1 and R2 must prefer the reachability to the RIP-learned prefixes via Serial link. If the Serial link is down, routers should prefer VLAN 60.

Solution:

All OSPF and Enhanced Interior Gateway Routing Protocol (EIGRP) prefixes redistributed into RIP will appear with the same metric on R2 showing the equal path to R1 via Serial link and VLAN 60.

By using an offset list on R2, you can change the metric on the prefixes received on VLAN 60 to a value that is less preferred than the metric over Serial link.

Use access list 0 to allow RIP to add the offset to all routes received on the interface:

R1:

```
offset-list 0 in 3 E0/1.60
```

R2:

```
offset-list 0 in 3 E0/0.60
```

R1:

```
R1#show ip route rip
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 148.49.160.6 to network 0.0.0.0

      148.49.0.0/16 is variably subnetted, 24 subnets, 2 masks
R        148.49.25.0/24 [120/1] via 148.49.21.2, 00:00:27, Serial1/0
R        148.49.50.0/24 [120/2] via 148.49.21.2, 00:00:27, Serial1/0
R        148.49.102.0/24 [120/1] via 148.49.21.2, 00:00:27, Serial1/0
R        148.49.105.0/24 [120/2] via 148.49.21.2, 00:00:27, Serial1/0
R1#
```

Note that R1 prefers the RIP networks via the Serial1/0 interface.

Note	To obtain a comprehensive view of the configuration tasks in this section, access the Mentor Guide engine. With the Mentor Guide engine, you can enter more than 1000 Cisco IOS Software commands, as well as a collection of proprietary commands such as **show all**.

5. IPv6 EIGRP

Issue: Configure IPv6 EIGRP AS 134 on the link FEC0::134:0/125 between routers R1, R3, and R4. R3 and R4 should not form the IPv6 EIGRP AS 134 neighbor relationship.

Solution:

Configure the IPv6 addresses provided in the “IPv6 IGP” diagram.

```
R1#show ipv6 interface brief e0/2
```



```
Ethernet0/2            [up/up]
    FE80::1
    FEC0::134:1
R1#

R3#show ipv6 interface brief e0/2
Ethernet0/2            [administratively down/down]
    FE80::3
    FEC0::134:3
R3#

R4#show ipv6 interface brief e0/2
Ethernet0/2            [up/up]
    FE80::4
    FEC0::134:4
R4#
```

Ensure that you can ping IPv6 between routers R1 and R3.

```
R1#ping FEC0::134:3
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to FEC0::134:3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/13 ms
R1#ping FEC0::134:4
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to FEC0::134:4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
R1#
```

Enable IPv6 unicast routing on the router with the **ipv6 unicast-routing** global configuration command. Configure the instances of IPv6 EIGRP AS 134 on the Ethernet02/ interfaces of R1, R3, and R4. According to the lab requirements, R3 and R4 should not form the IPv6 EIGRP AS 134 neighbor relationship. This requirement can be met by configuring the IPv6 EIGRP **neighbor** statements that point to the IPv6 link-local addresses.

R1:

```
interface Ethernet0/2
ipv6 eigrp 134
!
ipv6 router eigrp 134
 neighbor FE80::3 Ethernet0/2
 neighbor FE80::4 Ethernet0/2
!
```

R3:

```
interface Ethernet0/2
ipv6 eigrp 134
!
ipv6 router eigrp 134
 neighbor FE80::1 Ethernet0/2
```

R4:

```
interface Ethernet0/2
ipv6 eigrp 134
!
ipv6 router eigrp 134
 neighbor FE80::1 Ethernet0/2
```

Verify the IPv6 EIGRP neighbors on R1:

```
R1#show ipv6 eigrp neighbors
EIGRP-IPv6 Neighbors for AS(134)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q  Seq
                                                   (sec)         (ms)       Cnt Num
0   Link-local address:     Et0/2                    11 00:04:35 1279  5000  0  12
    FE80::3
1   Link-local address:     Et0/2                    10 01:40:25  106   636  0  11
    FE80::4
R1#
```

Note that R1 forms the IPv6 EIGRP neighbor relationships with R3 and R4.

Verify the IPv6 EIGRP neighbors on R3:

```
R3#show ipv6 eigrp neighbors
EIGRP-IPv6 Neighbors for AS(134)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q  Seq
                                                   (sec)         (ms)       Cnt Num
0   Link-local address:     Et0/2                    13 00:05:40    5   100  0  23
    FE80::1
R3#
```

Note that R3 forms the IPv6 EIGRP neighbor relationship only with R1.

Verify the IPv6 EIGRP neighbors on R4:

```
R4#show ipv6 eigrp neighbors
EIGRP-IPv6 Neighbors for AS(134)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q  Seq
                                                   (sec)         (ms)       Cnt Num
0   Link-local address:     Et0/2                    13 00:06:25    6   100  0  23
    FE80::1
R4#
```

Note that R4 forms the IPv6 EIGRP neighbor relationship only with R1.

Advertise the Loopback103 and the Loopback104 networks on routers R3 and R4. Disable the IPv6 EIGRP AS 134 split horizon on the Ethernet0/2 interface of R1 to allow route advertisements between R3 and R4.

R3:

```
interface Loopback103
ipv6 address FEC0::103:1/125
ipv6 eigrp 134
```

R4:

```
interface Loopback104
ipv6 address FEC0::104:1/125
 ipv6 eigrp 134
```

R1:

```
!
interface Ethernet0/2
no ipv6 split-horizon eigrp 134
```

Verify the IPv6 EIGRP networks in the IPv6 routing table. Here is an example from R4:

```
R4#show ipv6 route eigrp
IPv6 Routing Table - default - 9 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
```



```
       H - NHRP, I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
       IS - ISIS summary, D - EIGRP, EX - EIGRP external, NM - NEMO
       ND - ND Default, NDp - ND Prefix, DCE - Destination, NDr - Redirect
       O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, ls - LISP site
       ld - LISP dyn-EID, a - Application
D   FEC0::103:0/125 [90/435200]
     via FE80::1, Ethernet0/2
R4#
```

Note	To obtain a comprehensive view of the configuration tasks in this section, access the Mentor Guide engine. With the Mentor Guide engine, you can enter more than 1000 Cisco IOS Software commands, as well as a collection of proprietary commands such as **show all**.

6. IPv4 EIGRP

Issue: Configure EIGRP AS 200 between routers R1, R3, and R6 on links 148.49.160.0/24 and 148.49.63.0/24.
On R6, the interfaces Loopback106 and Loopback1 must be included in EIGRP AS 200.

Solution:

You can use the **network** statement configuration under the EIGRP AS 200 routing process with a wildcard mask to specifically include the desired links in the EIGRP topology:

R1:

```
router eigrp 200
 network 148.49.160.0 0.0.0.255
```

R3:

```
router eigrp 200
 network 148.49.63.0 0.0.0.255
```

R6:

```
router eigrp 200
 network 148.49.1.0 0.0.0.255
 network 148.49.63.0 0.0.0.255
 network 148.49.106.0 0.0.0.255
 network 148.49.160.0 0.0.0.255
```

Issue: Configure EIGRP AS 100 between R6 and R7.
Advertise the Loopback107 interface into EIGRP AS 100 from R7.

Solution:

You can use the **network** statement configuration under the EIGRP AS 100 routing process with a wildcard mask to specifically include the desired links in the EIGRP topology:

R6:

```
router eigrp 100
 network 148.49.76.0 0.0.0.255
```

R7:

```
router eigrp 100
 network 148.49.76.0 0.0.0.255
 network 148.49.107.0 0.0.0.255
```

Note	To obtain a comprehensive view of the configuration tasks in this section, access the Mentor Guide engine. With the Mentor Guide engine, you can enter more than 1000 Cisco IOS Software commands, as well as a collection of proprietary commands such as **show all**.

7. IPv4 Route Redistribution

Three interior gateway protocols (IGPs) are configured in this lab. Two of these IGPs, OSPF and EIGRP, are classified as core routing protocols; RIP is classified as an edge routing protocol. The core routing protocols provide transit services to the edge protocols. Core routing protocols can also provide a level of redundancy to the edge protocols. Three redistribution points will perform redistribution: R1, R3, and R6.

The redistribution strategy for this answer key is described below:

- On R6, EIGRP AS 100 prefixes are redistributed into EIGRP AS 200.
- On R6, EIGRP AS 200 prefixes are redistributed into EIGRP AS 100. This will inject RIP, OSPF, and EIGRP AS 200 prefixes into EIGRP AS 100, providing full IGP reachability to R7.
- EIGRP is redistributed into OSPF on routers R1 and R3. Only EIGRP-originated prefixes are permitted in this redistribution, to provide connectivity to networks of EIGRP AS 100 and AS 200 from the OSPF domain.
- OSPF is redistributed into EIGRP on routers R1 and R3. The EIGRP-originated prefixes are denied to stop possible route feedback, and all other prefixes are allowed.
- EIGRP is redistributed into RIP to supply reachability information to EIGRP-originated prefixes from the RIP domain.
- OSPF is redistributed into RIP to supply reachability information to OSPF-originated prefixes from the RIP domain.
- RIP prefixes are redistributed into OSPF, and EIGRP is performed on R1. Only RIP-originated prefixes are allowed in this redistribution instance.
- Change the administrative distance for the OSPF external prefixes representing the EIGRP-originated networks to something higher and less preferred than the administrative distance of the external EIGRP prefixes. Since the default values of EIGRP external routes are assigned the administrative distance of 170, the non-EIGRP-originated OSPF external routes need to have their administrative distance set to a value of 171.
- No redistribution control filtering is required between the EIGRP AS 100 and 200 instances on R6.
- The general practice applied to any **redistribute connected** configuration is to always apply a route map to select only the desired connected interfaces to be redistributed.
- On R6, use the **offset-list** EIGRP configuration command to make R6 prefer the Loopback108 interface network via R3. Apply the offset metric to the Loopback108 prefix that is learned via the VLAN 50 interface. This offset metric will make the Loopback108 prefix preferred via the VLAN 20 interface which is connected to R3.

- Add the **default-information originate** option on R4 to the NSSA configuration to provide a redundant path for R8 to the rest of the network if the Ethernet0/1 and Ethernet0/2 interfaces of R3 become unavailable. This redundant path enables you to verify the Network Address Translation (NAT) configuration in a later task.

Redistribution Table

The following table provides a useful summary of which prefixes were imported into a given routing protocol. Whenever a permit column for a given routing protocol is completely empty, it reflects that no prefixes were redistributed into the routing protocol. This means that the routing protocol is involved in one-way redistribution.

IPv4 Redistribution

Redist Point	Into RIP		Into OSPF		Into EIGRP	
	PERMIT	DENY	PERMIT	DENY	PERMIT	DENY
R1	OSPF native routes EIGRP native routes Connected	implicit	EIGRP native routes RIP native routes	implicit	Permit any	EIGRP native routes
R3			EIGRP native routes	implicit	Permit any	EIGRP native routes
R8			Connected			

A Tcl script can be used to test universal reachability. To use the script, enter the **tclsh** command in privileged mode and paste in this script. To stop failing pings, hold down the **Ctrl** and **Shift** keys while pressing the **6** key twice. After you are done, enter the **tclquit** command to leave Tcl mode.

Note Tcl connectivity verification scripts for each router are available in the "Verification" link in the web portal CIERSWB service tab.

```
tclsh
foreach address {
148.49.101.1
148.49.160.1
148.49.134.1
148.49.21.1
148.49.21.2
148.49.102.1
148.49.25.2
148.49.63.3
148.49.134.3
148.49.103.1
148.49.34.3
148.49.104.1
148.49.134.4
148.49.34.4
148.49.25.5
148.49.105.1
148.49.50.5
148.49.106.1
148.49.76.6
```

```
148.49.63.6
148.49.160.1
148.49.108.1
148.49.34.8
148.49.50.9
148.49.107.1
148.49.76.7} {ping $address}
```

You must ensure that the solution is stable. If there are split-horizon or other route feedback problems, routes may continually be inserted and removed from the routing tables. To test stability, observe the output of the **debug ip routing** command.

Issue: After all redistribution is performed, R9 still has not been included in any IGP and does not have **ip routing** enabled—but it can ping all IP addresses configured in this topology. How is R9 able to ping the rest of the network?

The answer is that a router that is configured with its default configuration will forward all IP packets out of an IP-configured interface and rely on an upstream locally attached router that has proxy Address Resolution Protocol (ARP) enabled. To prove this, ensure that the **no ip routing** command is performed on R9, then go to R5 and enable **debug arp**.

Verify the IP routing table on R9:

```
R9#show ip route
Default gateway is not set

Host               Gateway           Last Use    Total Uses  Interface
ICMP redirect cache is empty
R9#
```

Enable the **debug arp** command on R5:

```
R5#debug arp
ARP packet debugging is on
R5#
```

Once this is enabled, go back to R9 and ping any nonlocal IP address and examine the ARP messages generated on R5.

In the test below, the unassigned IP address of 140.10.1.1 is used:

```
R9#ping 140.10.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 140.10.1.1, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

As the **debug arp** output on R5 reflects, R9 sends the ARP requests for the nonlocal IP address 140.10.1.1:

```
R5#
*Aug  13  17:25:42.186:  IP  ARP:  rcvd  req  src  148.49.50.9  aabb.cc00.0d00,  dst
140.10.1.1 Ethernet0/0.40
R5#
*Aug  13  17:25:44.191:  IP  ARP:  rcvd  req  src  148.49.50.9  aabb.cc00.0d00,  dst
140.10.1.1 Ethernet0/0.40
R5#
```

Note that R5 does not reply to the ARP request because R5 does not have a route to 140.10.1.1.

Now ping one of the IP addresses that is configured in this lab (for example, the Loopback101 IP address 148.49.101.1 from R9) and verify the output of the **debug arp** command on R5:

```
R9#ping 148.49.101.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 148.49.101.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 7/209/1013 ms
R9#

R5#
*Aug  13  17:30:57.255:  IP  ARP:  rcvd  req  src  148.49.50.9  aabb.cc00.0d00,  dst
148.49.101.1 Ethernet0/0.40
*Aug 13 17:30:57.255: IP ARP: sent rep src 148.49.101.1 aabb.cc00.0500,
                 dst 148.49.50.9 aabb.cc00.0d00 Ethernet0/0.40
R5#
```

Note that R5 received the ARP request for the IP address 140.10.1.1 and replied with its own MAC address.

R9 added the ARP resolution in the table:

```
R9#show arp
Protocol  Address          Age (min)  Hardware Addr   Type   Interface
Internet  148.49.50.9              -   aabb.cc00.0d00  ARPA   Ethernet0/0
Internet  148.49.101.1             3   aabb.cc00.0500  ARPA   Ethernet0/0
R9#
```

This clearly proves that a router with the disabled routing functionality will generate ARPs for non-locally addressed IP packets with the expectation that a locally attached router will respond to the ARP request due to proxy ARP. Therefore, with proxy ARP enabled on a locally attached router, a router such as R9 in this scenario will be able to reach all destinations without any IGP or static configuration.

This approach meets the requirements of this lab, which forbid any IGP or default gateway to be configured in R9.

Note	To obtain a comprehensive view of the configuration tasks in this section, access the Mentor Guide engine. With the Mentor Guide engine, you can enter more than 1000 Cisco IOS Software commands, as well as a collection of proprietary commands such as **show all**.

8. IPv6 Route Redistribution

Issue: Perform mutual redistribution between IPv6 EIGRP and OSPFv3.

Solution:

Mutual redistribution is required to provide connectivity between all the IPv6 networks advertised in the OSPFv3 and IPv6 EIGRP domains.

Mutually redistribute OSPFv3 and IPv6 EIGRP on R1:

```
ipv6 router eigrp 134
 neighbor FE80::3 Ethernet0/2
 neighbor FE80::4 Ethernet0/2
 redistribute ospf 1
 default-metric 1000 100 3 255 1500
```

```
ipv6 router ospf 1
 redistribute eigrp 134
```

Connected IPv6 networks will not be redistributed automatically from OSPFv3 into IPv6 EIGRP and vice versa.

Use the **redistribute connected** command to redistribute the Loopback101 network into IPv6 EIGRP and the Ethernet0/2 network into OSPFv3. To ensure that you do not redistribute other connected networks, configure a route map to precisely limit the number of connected networks that will be redistributed.

R1:

```
ipv6 router eigrp 134
 neighbor FE80::3 Ethernet0/2
 neighbor FE80::4 Ethernet0/2
 redistribute connected route-map Conn-->EIGRP_V6
 redistribute ospf 1
 default-metric 1000 100 3 255 1500
!
ipv6 router ospf 1
 redistribute connected metric 10 route-map Conn-->OSPF_V6
 redistribute eigrp 134
!
ipv6 prefix-list LOOP101IPV6 seq 5 permit FEC0::101:0/125
ipv6 prefix-list LOOP101IPV6 seq 10 permit FEC0::160:0/125
!
ipv6 prefix-list NET-134 seq 5 permit FEC0::134:0/125
!
route-map Conn-->OSPF_V6 permit 10
 match ipv6 address prefix-list NET-134
!
!
route-map Conn-->EIGRP_V6 permit 10
 match interface Loopback101 Ethernet0/1.50
!
!
```

This script can be typed once and pasted into each router to test connectivity:

```
tclsh
foreach address {
fec0::101:1
fec0::160:1
fec0::134:1
fec0::103:1
fec0::134:3
fec0::104:1
fec0::134:4
fec0::106:1
fec0::160:6} {ping $address}
```

Note	To obtain a comprehensive view of the configuration tasks in this section, access the Mentor Guide engine. With the Mentor Guide engine, you can enter more than 1000 Cisco IOS Software commands, as well as a collection of proprietary commands such as **show all**.

9. Border Gateway Protocol

Issue: Do not form a full mesh of IBGP peer relationships within AS 314. R8 will be used to exchange the BGP updates between R3 and R4.

Solution:

R3, R4, and R8 are all in AS 314. Therefore, these IBGP speakers need to be peered to each other according to the BGP RFC, because IBGP peers will not forward IBGP-learned updates to another IBGP peer. This rule is relaxed with a route reflector configuration. The configuration task directs you to configure R8 to forward IBGP updates between R3 and R4. Therefore, R8 will be the route reflector and R3 and R4 the route reflector clients.

Issue: Use the Loopback interface 106 and 104 for the peer relationship between R6 and R4.

Solution:

When you peer R6 and R4, remember to use the **ebgp-multihop** keyword on the **neighbor** statements to increase the default TTL value of 1 for the EBGP TCP/IP session. The maximum value of 255 is used in this answer key. Also configure the **BGP update-source** command to terminate the EBGP sessions between the Loopback interfaces 104 and 106. If you fail to do so, BGP will attempt to initiate the EBGP session from the outgoing interface.

Issue: Advertise the networks associated to the following loopback interfaces in AS 314 with the ORIGIN attribute IGP: Loopbacks 3333, 4444, and 8888.

Solution:

Because all these networks are /32 networks, ensure that you specify a precise mask in the **network** statement under the BGP process (for example, **network 3.3.3.3 mask 255.255.255.255**).

You must use the BGP network statement because of the configuration requirement to originate the prefixes into BGP as IGP. When you originate a prefix into BGP with the **network** configuration command, the prefix is assigned with IGP.

Disable synchronization, or the route reflector will not reflect the prefixes that are not in the routing table.

Provided below is the output on one of the routers in AS 314. The other routers will have a similar output. Check the Mentor Guide engine for more details.

```
R4#sh ip bgp summary | begin ^Neig
Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
148.49.34.8     4   314     396     394        8    0    0 01:40:29        2
148.49.106.1    4    67       6       6        8    0    0 00:00:20        0

R4#sh ip bgp
BGP table version is 8, local router ID is 148.49.104.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*>i3.3.3.3/32       148.49.34.3              0    100      0 i
*> 4.4.4.4/32       0.0.0.0                  0         32768 i
*>i8.8.8.8/32       148.49.34.8              0    100      0 i
```

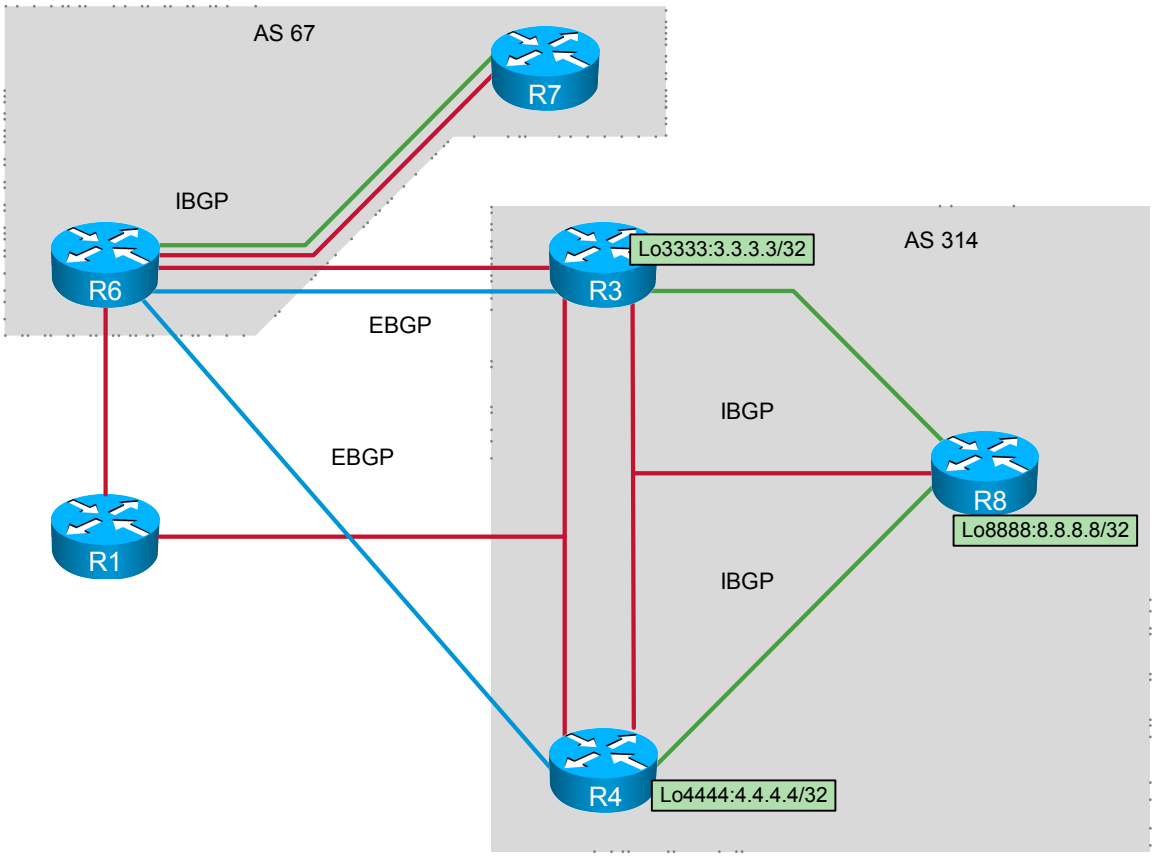

Note	To obtain a comprehensive view of the configuration tasks in this section, access the Mentor Guide engine. With the Mentor Guide engine, you can enter more than 1000 Cisco IOS Software commands, as well as a collection of proprietary commands such as **show all**.

10. Quality of Service

Issue: HTTP traffic between routers R1 and R2 must be discarded without any further system processing. Use the Modular QoS CLI (MQC) to accomplish this task.

Solution:

The Cisco IOS Software feature Modular QoS CLI Unconditional Packet Discard allows you to accomplish this task. The key is "without any further system processing," meaning fast, unconditional discard.

Classify the traffic you want to discard unconditionally on both routers R1 and R2:

```
class-map match-all HTTPMAP
 match access-group name HTTPACL
!
ip access-list extended HTTPACL
 permit tcp any any eq www
```


Create a policy map for the action you want to perform on the classified traffic:

```
policy-map HTTPPOLICY
 class HTTPMAP
  drop
```

Apply your policy on the Serial and VLAN 60 interfaces between R1 and R2, on both routers in the output direction:

```
interface Serial1/0
service-policy output HTTPPOLICY
!
int E0/0.60
service-policy output HTTPPOLICY
```

To verify, enter the following commands on R1:

```
R1#telnet 148.49.21.2 www
Trying 148.49.21.2, 80 ...
% Connection timed out; remote host not responding

R1#show policy-map interface s1/0
 Serial1/0

  Service-policy output: HTTPPOLICY

    Class-map: HTTPMAP (match-all)
      1 packets, 48 bytes
      5 minute offered rate 0000 bps, drop rate 0000 bps
      Match: access-group name HTTPACL
      drop

    Class-map: class-default (match-any)
      1911 packets, 240945 bytes
      5 minute offered rate 0000 bps, drop rate 0000 bps
      Match: any
R1#
R1#show access-lists HTTPACL
Extended IP access list HTTPACL
    10 permit tcp any any eq www (1 match)
R1#
```

Note that the HTTP traffic is dropped on the Serial link according to the traffic classification criteria of this lab.

Note	To obtain a comprehensive view of the configuration tasks in this section, access the Mentor Guide engine. With the Mentor Guide engine, you can enter more than 1000 Cisco IOS Software commands, as well as a collection of proprietary commands such as **show all**.

11. System Administration

Issue: On R5, configure the username **admin** and password **lab**.
Configure system logging for both successful and failed login attempts.

Solution:

The Cisco IOS Software feature Cisco IOS Login Enhancements will help you to configure this task and other tasks from this section.

```
username admin password 0 lab
!
login on-success log
```

Issue: Provide a solution on R5 to slow down dictionary attacks that attempt to gain access to the username and access information. Introduce a 3-second delay between successive login attempts. If two login attempts to R5 fail within 15 seconds, R5 should not allow any login attempts for 10 seconds, except for the attempts coming from 148.49.25.2.

Solution:

A Cisco IOS Software device can accept virtual connections as fast as they can be processed.

Introducing a delay between login attempts helps to protect your router from a possible dictionary attack, which attempts to gain access to your username and password information. Delays can be enabled with the following command:

```
login delay 3
```

If the configured number of connection attempts fails within a specified time period, the Cisco IOS Software device will not accept any additional connections for a "quiet period." The number of failed connection attempts that trigger the quiet period can be specified via the new global configuration mode command **login block-for**. The predefined ACL that is excluded from the quiet period can be specified via the new global configuration mode command **login quiet-mode access-class**.

```
login block-for 10 attempts 2 within 15
login quiet-mode access-class QUIET-ACL
!
ip access-list standard QUIET-ACL
 permit 148.49.25.2
!
line vty 0 4
login local
```

Use Telnet from R9 to R5 and provide the correct credentials:

```
R9#telnet 148.49.50.5
Trying 148.49.50.5 ... Open

------------------------------------------------
Cisco 360 R&S Exercise Workbook
Product, POD location:  cierswbv5-ce-lab09-sc, SJ
Device:                 R5
------------------------------------------------


User Access Verification

Username: admin
```



```
Password:
R5>exit

[Connection to 148.49.50.5 closed by foreign host]
R9#
```

Verify logging on R5:

```
R5#
*Aug  13  21:03:52.364:  %SEC_LOGIN-5-LOGIN_SUCCESS:  Login  Success  [user:  admin]
[Source: UNKNOWN] [localport: 23] at 13:03:52 PST Tue Aug 13 2013
R5#
```

Use Telnet from R9 to R5 and make two failed login attempts:

```
R9#telnet 148.49.50.5
Trying 148.49.50.5 ... Open

-------------------------------------------------
Cisco 360 R&S Exercise Workbook
Product, POD location:  cierswbv5-ce-lab09-sc, SJ
Device:                 R5
-------------------------------------------------


User Access Verification

Username: aaa
Password:
% Login invalid

Username: aaa
Password:
% Login invalid

[Connection to 148.49.50.5 closed by foreign host]
```

Verify logging on R5:

```
R5#
*Aug  13  21:04:21.301:  %SEC_LOGIN-1-QUIET_MODE_ON:  Still  timeleft  for  watching
failures is 6 secs, [user: aaa] [Source: UNKNOWN] [localport: 23] [Reason: Login
Authentication Failed - BadUser] [ACL: QUIET-ACL] at 13:04:21 PST Tue Aug 13 2013
R5#
*Aug 13 21:04:31.305: %SEC_LOGIN-5-QUIET_MODE_OFF: Quiet Mode is OFF, because block
period timed out at 13:04:31 PST Tue Aug 13 2013
R5#
```

The router stays in quiet mode for 10 seconds:

```
R9#telnet 148.49.50.5
Trying 148.49.50.5 ...
% Connection refused by remote host

R9#telnet 148.49.50.5
Trying 148.49.50.5 ...
% Connection refused by remote host

R9#
```

In 10 seconds, the Telnet login prompt becomes available again:

```
R9#telnet 148.49.50.5
```

```
Trying 148.49.50.5 ... Open

------------------------------------------------
Cisco 360 R&S Exercise Workbook
Product, POD location:  cierswbv5-ce-lab09-sc, SJ
Device:                 R5
------------------------------------------------


User Access Verification

Username:
```

Note	To obtain a comprehensive view of the configuration tasks in this section, access the Mentor Guide engine. With the Mentor Guide engine, you can enter more than 1000 Cisco IOS Software commands, as well as a collection of proprietary commands such as **show all**.

12. Address Administration

Issue: Two groups consisting of virtual workstations are connected to VLAN 30. One HSRP group (GROUP1) is configured for the 148.49.34.1 gateway and the other HSRP group (GROUP2) is configured for 148.49.34.254.

Solution:

You need to configure Multigroup HSRP (MHSRP).

Essentially, R3 is configured with two HSRP groups (for example, GROUP1 and GROUP2) and R4 is configured with the same HSRP groups.

- For GROUP1, R3 is the active router and R4 is the standby router.
- For GROUP2, R4 is the active router and R3 is the standby router.

As a result of this configuration, workstations of GROUP1 will use R3 as a default gateway because the virtual IP address 148.49 34.1 is assigned to standby group 1. Likewise, workstations in GROUP2 will use R4 as a default gateway because the virtual IP address 148.49 34.254 is assigned to standby group 2.

Issue: GROUP1 should use router R3 as a preferred gateway. If the Ethernet0/2 interface fails on R3, R4 must be preferred for the duration of the failure.
GROUP2 should use router R4 as a preferred gateway. If the Ethernet0/2 interface fails on R4, R3 must be preferred for the duration of the failure.

Solution:

Interface tracking needs to be configured for the HSRP groups. The **standby track** command allows you to specify another interface on the router for the HSRP process to monitor in order to alter the HSRP priority for a given group. If the line protocol of the specified interface goes down, the HSRP priority is reduced. This means that another HSRP router with higher priority can become the active router if it has **standby preempt** enabled.

R3	R4
`interface Ethernet0/0` `encapsulation dot1Q 30`	`interface Ethernet0/0` `ip address 148.49.34.4`


```
 ip address 148.49.34.3 255.255.255.0
 standby 1 track Ethernet0/2
 standby 1 ip 148.49.34.1
 standby 1 priority 105
 standby 1 preempt
 standby 2 ip 148.49.34.254
 standby 2 preempt
```

```
 255.255.255.0
 standby 1 ip 148.49.34.1
 standby 1 preempt
 standby 2 ip 148.49.34.254
 standby 2 priority 105
 standby 2 preempt
 standby 2 track Ethernet0/2
```

Verify the HSRP configuration on R3:

```
R3#sh standby brief
                     P indicates configured to preempt.
                     |
Interface Grp  Pri P State   Active          Standby         Virtual IP
E0/0      1    105 P Active  local           148.49.34.4     148.49.34.1
E0/0      2    100 P Standby 148.49.34.4     local           148.49.34.254


R3#sh standby
Ethernet0/0 - Group 1
  State is Active
    7 state changes, last state change 01:47:32
  Virtual IP address is 148.49.34.1
  Active virtual MAC address is 0000.0c07.ac01
    Local virtual MAC address is 0000.0c07.ac01 (v1 default)
  Hello time 3 sec, hold time 10 sec
    Next hello sent in 0.672 secs
  Preemption enabled
  Active router is local
  Standby router is 148.49.34.4, priority 100 (expires in 10.640 sec)
  Priority 105 (configured 105)
    Track object 1 state Up decrement 10
  Group name is "GROUP1" (cfgd)
Ethernet0/0 - Group 2
  State is Standby
    3 state changes, last state change 03:13:34
  Virtual IP address is 148.49.34.254
  Active virtual MAC address is 0000.0c07.ac02
    Local virtual MAC address is 0000.0c07.ac02 (v1 default)
  Hello time 3 sec, hold time 10 sec
    Next hello sent in 1.456 secs
  Preemption enabled
  Active router is 148.49.34.4, priority 105 (expires in 10.400 sec)
  Standby router is local
  Priority 100 (default 100)
  Group name is "GROUP2" (cfgd)
R3#
R3#show track
Track 1
  Interface Ethernet0/2 line-protocol
  Line protocol is Up
    3 changes, last change 01:51:38
  Tracked by:
    HSRP Ethernet0/0 1
R3#
```

Verify the HSRP configuration on R4:

```
R4#sh standby brief
                     P indicates configured to preempt.
                     |
Interface   Grp  Pri P State   Active          Standby         Virtual IP
Et0/0       1    100 P Standby 148.49.34.3     local           148.49.34.1
```

```
Et0/0       2    105 P Active  local           148.49.34.3     148.49.34.254
R4#
R4#show standby
Ethernet0/0 - Group 1
  State is Standby
    7 state changes, last state change 01:49:59
  Virtual IP address is 148.49.34.1
  Active virtual MAC address is 0000.0c07.ac01
    Local virtual MAC address is 0000.0c07.ac01 (v1 default)
  Hello time 3 sec, hold time 10 sec
    Next hello sent in 2.384 secs
  Preemption enabled
  Active router is 148.49.34.3, priority 105 (expires in 8.896 sec)
  Standby router is local
  Priority 100 (default 100)
  Group name is "GROUP1" (cfgd)
Ethernet0/0 - Group 2
  State is Active
    2 state changes, last state change 03:25:22
  Virtual IP address is 148.49.34.254
  Active virtual MAC address is 0000.0c07.ac02
    Local virtual MAC address is 0000.0c07.ac02 (v1 default)
  Hello time 3 sec, hold time 10 sec
    Next hello sent in 2.464 secs
  Preemption enabled
  Active router is local
  Standby router is 148.49.34.3, priority 100 (expires in 9.600 sec)
  Priority 105 (configured 105)
    Track object 1 state Up decrement 10
  Group name is "GROUP2" (cfgd)
R4# R4#show track
Track 1
  Interface Ethernet0/2 line-protocol
  Line protocol is Up
    1 change, last change 03:27:33
  Tracked by:
    HSRP Ethernet0/0 2
R4#
```

Issue: The Telnet session from the Loopback108 interface of R8 to the Loopback106 interface of R6 should be listed on router R6 as originating from the IP address 148.49.108.180.

Solution:

This task requires that you rewrite the source IP address of the Telnet session from 148.49.108.1 to 148.49.108.180.

NAT is an obvious choice to use. However, there are other configuration requirements, which are listed below.

Issue: This Telnet session must be established via R3.
If the path via R3 becomes unavailable because the Ethernet0/1 and Ethernet0/2 interfaces of R3 have failed, the Telnet session should be established via R4. Extend your solution to the IP Finger protocol as well.

Solution:

Note that this solution interacts with the OSPF NSSA requirement described in the OSPF section. The preferred path from R8 to NSSA network is via R3, because the EIGRP AS 200 prefixes are redistributed into the OSPF NSSA Area 30 on R3. If the Ethernet0/1 and Ethernet0/2

interfaces of R3 fail, R8 should be able to forward to the Loopback106 network via R4, because R4 originates the default network into the OSPF NSSA Area 30.

Configure match criteria for the traffic to be sent to the NAT process on routers R3 and R4:

```
ip access-list extended SOURCE108
 permit tcp host 148.49.108.1 host 148.49.106.1 eq telnet
 permit tcp host 148.49.108.1 host 148.49.106.1 eq finger
!
route-map MAP108 permit 10
 match ip address SOURCE108
```

Configure a NAT pool for the translation for the Telnet session on routers R3 and R4:

```
ip nat pool TELNET180 148.49.108.180 148.49.108.180 prefix-length 24
```

Configure the network address translation rule by applying the criteria specified in the previous two steps:

```
ip nat inside source route-map MAP108 pool TELNET180
```

Configure NAT interfaces on R3 and R4.

R3:

```
interface Ethernet0/0
 ip nat inside
interface Ethernet0/1
 ip nat outside
interface Ethernet0/2
 ip nat outside
```

R4:

```
interface Ethernet0/0
 ip nat inside
interface Ethernet0/2
 ip nat outside
```

Configure NAT configuration on R3 and R4.

```
R3#sh ip nat statistics
Total active translations: 0 (0 static, 0 dynamic; 0 extended)
Peak translations: 1, occurred 00:26:14 ago
Outside interfaces:
  Ethernet0/1, Ethernet0/2
Inside interfaces:
  Ethernet0/0
Hits: 155  Misses: 0
CEF Translated packets: 155, CEF Punted packets: 0
Expired translations: 2
Dynamic mappings:
-- Inside Source
[Id: 1] route-map MAP108 pool TELNET180 refcount 0
 pool TELNET180: netmask 255.255.255.0
        start 148.49.108.180 end 148.49.108.180
        type generic, total addresses 1, allocated 0 (0%), misses 0

Total doors: 0
Appl doors: 0
Normal doors: 0
```

```
Queued Packets: 0
R3#
R4#sh ip nat statistics
Total active translations: 0 (0 static, 0 dynamic; 0 extended)
Peak translations: 1, occurred 00:28:04 ago
Outside interfaces:
  Ethernet0/2
Inside interfaces:
  Ethernet0/0
Hits: 95  Misses: 0
CEF Translated packets: 95, CEF Punted packets: 0
Expired translations: 1
Dynamic mappings:
-- Inside Source
[Id: 1] route-map MAP108 pool TELNET180 refcount 0
 pool TELNET180: netmask 255.255.255.0
        start 148.49.108.180 end 148.49.108.180
        type generic, total addresses 1, allocated 0 (0%), misses 0

Total doors: 0
Appl doors: 0
Normal doors: 0
Queued Packets: 0
R4#
```

To test your configuration, perform the following steps:

Step 1 Open the Telnet session from R8:

```
R8#telnet 148.49.106.1 /source-interface loopback 108
Trying 148.49.106.1 ... Open

------------------------------------------------
Cisco 360 R&S Exercise Workbook
Product, POD location:  cierswbv5-ce-lab09-sc, SJ
Device:                 R6
------------------------------------------------
R6#who
    Line       User       Host(s)              Idle       Location
   0 con 0                 idle                 00:00:41   cierswbv5-ce-lab09-sc, SJ
*  2 vty 0                 idle                 00:00:00 148.49.108.180

  Interface    User               Mode         Idle     Peer Address

R6#exit

[Connection to 148.49.106.1 closed by foreign host]
R8#
```

Note that the **who** command output shows the translated IP address 148.49.108.180 for the Telnet session.

Step 2 R8 should have R3 as the preferred next hop, according to IGP configuration and redistribution. Check NAT translation on both R3 and R4 for this Telnet session:

```
R3#show ip nat translations
Pro Inside global      Inside local       Outside local      Outside global
tcp 148.49.108.180:26571 148.49.108.1:26571 148.49.106.1:23  148.49.106.1:23
R3#

R4#show ip nat translations
R4#
```

Note that R3 performs the NAT operations.

Step 3 Shut down the Ethernet0/1 and Ethernet0/2 interfaces on R3. R8 should lose the more specific route to 148.49.106.0 from R3 and start using 0.0.0.0 originated from R4.

```
R3#clear ip nat translation *
R3#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R3(config)#int e0/1
R3(config-if)#shut
R3(config-if)#int e0/2
R3(config-if)#shut
R3(config-if)#
```

Step 4 Open the Telnet session from R8:

```
R8#telnet 148.49.106.1 /source-interface loopback 108
Trying 148.49.106.1 ... Open

------------------------------------------------
Cisco 360 R&S Exercise Workbook
Product, POD location:  cierswbv5-ce-lab09-sc, SJ
Device:                 R6
------------------------------------------------
R6#who
    Line       User       Host(s)              Idle       Location
   0 con 0                 idle                 00:00:41   cierswbv5-ce-lab09-sc, SJ
*  2 vty 0                 idle                 00:00:00 148.49.108.180

  Interface    User               Mode         Idle     Peer Address

R6#exit

[Connection to 148.49.106.1 closed by foreign host]
R8#
```

Step 5 Check NAT translation on R4 for this Telnet session:

```
R4#show ip nat translations
Pro Inside global      Inside local       Outside local      Outside global
tcp 148.49.108.180:56613 148.49.108.1:56613 148.49.106.1:23  148.49.106.1:23
R4#
```

Note that when the Ethernet0/1 and Ethernet0/2 interface are shut down on R3, R4 performs the NAT operations.

Step 6 Enable the Ethernet0/1 and Ethernet0/2 interfaces on R3.

```
R3(config-if)#no shut
R3(config-if)#int e0/1
R3(config-if)#no shut
R3(config-if)#end
R3#
```

Note	To obtain a comprehensive view of the configuration tasks in this section, access the Mentor Guide engine. With the Mentor Guide engine, you can enter more than 1000 Cisco IOS Software commands, as well as a collection of proprietary commands such as **show all**.

13. Multicast

Issue: Configure Loopback interface 1 on router R4 with the IP address 148.49.1.1/32. Configure Loopback interface 1 on router R6 with the IP address 148.49.1.1/32.

Solution:

Configuring the same IP addresses on two different routers may seem confusing. In the unicast world, applications seldom involve exact duplicate IP addresses assigned to two different devices; however, in the multicast world, this technique is used when configuring anycasting for sparse-mode rendezvous points (RPs).

Issue: Use these IP addresses for the RPs on R4 and R6 as well as mapping agents on the same routers.

Solution:

RPs and mapping agents direct you to configure PIM Auto-Discovery to fulfill the group-to-RP mapping requirements for sparse mode. Only PIM Auto-Discovery uses a mapping agent. Both R4 and R6 will be configured as RPs and mapping agents using the same IP unicast address. The RPs will be peered to each other with the Multicast Source Discovery Protocol (MSDP) peer relationship to exchange the IP addresses of multicast active sources that have registered with either RP. With MSDP, the messages used by MSDP peers to advertise multicast active sources are called source active messages. Examine the following diagram displaying the multicast network and MSDP peer relationship.

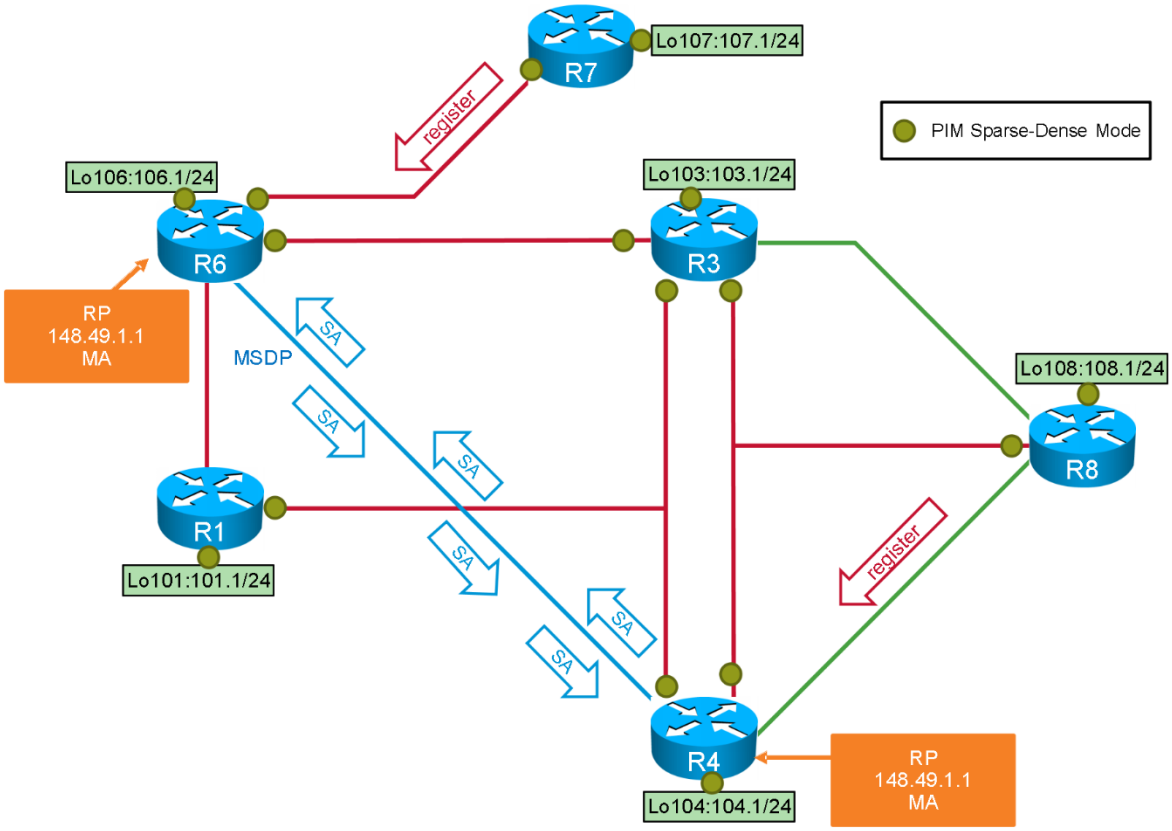

If multicast traffic is sourced to 225.23.23.23 from R8, R8 will send its registration message to the closest RP that is defined by the lookup in the routing table for the entry 148.49.1.1/32, which would be R4:

```
R8#sh ip route | inc 148.49.1.1
O IA    148.49.1.1/32 [110/2] via 148.49.34.4, 02:22:09, Vlan30
```

R8 will send the registration message to R4, and R4 will send the source address of R8 to R6 via MSDP peer.

Here is the configuration:

R1:

```
ip multicast-routing
!
interface Loopback101
 ip address 148.49.101.1 255.255.255.0
 ip pim sparse-dense-mode
!
interface Ethernet0/1.50
 encapsulation dot1Q 50
 ip address 148.49.160.1 255.255.255.0
ip pim sparse-dense-mode
!
interface Ethernet0/2
```

```
 ip address 148.49.134.1 255.255.255.0
 ip pim sparse-dense-mode
```

R3:

```
ip multicast-routing
!
interface Loopback103
 ip address 148.49.103.1 255.255.255.0
 ip pim sparse-dense-mode
 ip igmp join-group 225.23.23.23
!
interface Ethernet0/0
 ip address 148.49.34.3 255.255.255.0
 ip pim sparse-dense-mode

interface Ethernet0/2
 ip address 148.49.134.3 255.255.255.0
 ip pim sparse-dense-mode
```

R4:

```
ip multicast-routing
!
interface Loopback104
 ip address 148.49.104.1 255.255.255.0
 ip pim sparse-dense-mode
 ip igmp join-group 225.23.23.23
! interface Ethernet0/0
 ip address 148.49.34.4 255.255.255.0
 ip pim sparse-dense-mode
!
interface Ethernet0/2
 ip address 148.49.134.4 255.255.255.0
 ip pim sparse-dense-mode
!
ip pim accept-rp auto-rp
ip pim send-rp-announce Loopback1 scope 10 interval 20
ip pim send-rp-discovery Loopback1 scope 10
ip msdp peer 148.49.106.1 connect-source Loopback104
!
```

R6:

```
ip multicast-routing
interface Loopback106
 ip address 148.49.106.1 255.255.255.0
 ip pim sparse-dense-mode
 ip igmp join-group 225.23.23.23
!
interface Ethernet0/0.10
 encapsulation dot1Q 10
 ip address 148.49.76.6 255.255.255.0
ip pim sparse-dense-mode
!
interface Ethernet0/0.20
 encapsulation dot1Q 20
 ip address 148.49.63.6 255.255.255.0
 no ip redirects
 ip pim sparse-dense-mode
!
interface Ethernet0/0.50
 encapsulation dot1Q 50
 ip address 148.49.160.6 255.255.255.0
```



```
 no ip redirects
 ip pim sparse-dense-mode
! ip pim accept-rp auto-rp
ip pim send-rp-announce Loopback1 scope 10 interval 20
ip pim send-rp-discovery Loopback1 scope 10
ip msdp peer 148.49.104.1 connect-source Loopback106
ip msdp originator-id Loopback106
```

R7:

```
ip multicast-routing
!
interface Loopback107
 ip address 148.49.107.1 255.255.255.0
 ip pim sparse-dense-mode
 ip igmp join-group 225.23.23.23
!
interface Ethernet0/0
 ip address 148.49.76.7 255.255.255.0
 ip pim sparse-dense-mode
!
```

R8:

```
ip multicast-routing
!
interface Loopback108
 ip address 148.49.108.1 255.255.255.0
 ip pim sparse-dense-mode
 ip igmp join-group 225.23.23.23
!
interface Ethernet0/0
 ip address 148.49.34.8 255.255.255.0
 ip pim sparse-dense-mode
 ip ospf priority 0
!
```

Verification:

Ensure that all multicast routers learn the RP:

```
R8#show ip pim rp map
PIM Group-to-RP Mappings

Group(s) 224.0.0.0/4
  RP 148.49.1.1 (?), v2v1
    Info source: 148.49.1.1 (?), elected via Auto-RP
         Uptime: 1d16h, expires: 00:02:15
R8#

R4#show ip pim rp map
PIM Group-to-RP Mappings
This system is an RP (Auto-RP)
This system is an RP-mapping agent (Loopback1)

Group(s) 224.0.0.0/4
  RP 148.49.1.1 (?), v2v1
    Info source: 148.49.1.1 (?), elected via Auto-RP
         Uptime: 1d16h, expires: 00:00:54
R4#

R6#show ip pim rp map
PIM Group-to-RP Mappings
This system is an RP (Auto-RP)
This system is an RP-mapping agent (Loopback1)
```

```
Group(s) 224.0.0.0/4
  RP 148.49.1.1 (?), v2v1
    Info source: 148.49.1.1 (?), elected via Auto-RP
         Uptime: 1d16h, expires: 00:00:57
R6#

R7#show ip pim rp map
PIM Group-to-RP Mappings

Group(s) 224.0.0.0/4
  RP 148.49.1.1 (?), v2v1
    Info source: 148.49.1.1 (?), elected via Auto-RP
         Uptime: 1d00h, expires: 00:02:16
R7#

R3#show ip pim rp map
PIM Group-to-RP Mappings

Group(s) 224.0.0.0/4
  RP 148.49.1.1 (?), v2v1
    Info source: 148.49.1.1 (?), elected via Auto-RP
         Uptime: 1d00h, expires: 00:02:01
R3#

R1#show ip pim rp map
PIM Group-to-RP Mappings

Group(s) 224.0.0.0/4
  RP 148.49.1.1 (?), v2v1
    Info source: 148.49.1.1 (?), elected via Auto-RP
         Uptime: 1d16h, expires: 00:02:45
R1#
```

Verify the MSDP peer relationship:

```
R4#show ip msdp summary
MSDP Peer Status Summary
Peer Address     AS    State    Uptime/  Reset SA    Peer Name
                                Downtime Count Count
148.49.106.1     67    Up       1d00h    0     0     ?
R4#

R6#show ip msdp summary
MSDP Peer Status Summary
Peer Address     AS    State    Uptime/  Reset SA    Peer Name
                                Downtime Count Count
148.49.104.1     314   Up       1d00h    0     0     ?
R6#
```

Ping 225.23.23.23, source from the Ethernet0/0 interfaces of R8 and R7, respectively:

```
R8#ping 225.23.23.23 source 148.49.34.8
Type escape sequence to abort.
Sending 1, 100-byte ICMP Echos to 225.23.23.23, timeout is 2 seconds:
Packet sent with a source address of 148.49.34.8

Reply to request 0 from 148.49.104.1, 1 ms
Reply to request 0 from 148.49.107.1, 1 ms
Reply to request 0 from 148.49.106.1, 1 ms
Reply to request 0 from 148.49.101.1, 1 ms
Reply to request 0 from 148.49.108.1, 1 ms
Reply to request 0 from 148.49.103.1, 1 ms
R8#
```

Note that you may have to ping from R8 a few times to receive all necessary replies.


```
R6#sh ip msdp sa-cache
MSDP Source-Active Cache - 1 entries
(148.49.34.8, 225.23.23.23), RP 148.49.104.1, BGP/AS 0, 00:00:39/00:05:20, Peer
148.49.104.1
R6#
```

Similarly, if you ping the group 225.23.23.23 from R7, R7 will send its registration message to R6, and R6 will send a source active to R4:

```
R7#ping 225.23.23.23 source Ethernet0/0
Type escape sequence to abort.
Sending 1, 100-byte ICMP Echos to 225.23.23.23, timeout is 2 seconds:
Packet sent with a source address of 148.49.76.7

Reply to request 0 from 148.49.107.1, 1 ms
Reply to request 0 from 148.49.108.1, 2 ms
Reply to request 0 from 148.49.104.1, 2 ms
Reply to request 0 from 148.49.101.1, 1 ms
Reply to request 0 from 148.49.103.1, 1 ms
Reply to request 0 from 148.49.106.1, 1 ms
R7#

R4#sh ip msdp sa-cache
MSDP Source-Active Cache - 1 entries
(148.49.76.7, 225.23.23.23), RP 148.49.106.1, BGP/AS 0, 00:00:35/00:05:55, Peer
148.49.106.1
R4#
```

Issue: Ensure that R4 and R6 accept join and prune messages only for RPs that are in their Auto-RP cache.

Solution:

To configure a router to accept join or prune messages destined for a specified RP and for a specific list of groups, use the **ip pim accept-rp** command in global configuration mode.

R4 and R6:

```
ip pim accept-rp auto-rp
```

Issue: All mapping agents should send RP announcements three times faster than the default.

Solution:

By default, a mapping agent sends RP announcements every 60 seconds. Change the interval to 20.

R4 and R6:

```
ip pim send-rp-announce Loopback 1 scope 10 interval 20
```

Note	To obtain a comprehensive view of the configuration tasks in this section, access the Mentor Guide engine. With the Mentor Guide engine, you can enter more than 1000 Cisco IOS Software commands, as well as a collection of proprietary commands such as **show all**.

