360RSW05-WKB-CA02

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 2 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 2 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: Intermediate to Advanced

Restrictions and Goals

Note	Read this section carefully.

- To receive credit for a subsection, you must fully complete the subsection per the requirements. You will *not* receive partial credit for partially completed subsections.
- IPv4 subnets that are displayed in the scenario diagram belong to network 172.16.0.0/16.
- *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**).
- Do not use the **ip default-network** command.
- All IP addresses from the 172.16.0.0/16 range that are involved in this scenario must be reachable, unless explicitly specified otherwise.

- The OSPF PID 1111 network is initialized for the DMVPN configuration. Subnets that are routing via OSPF PID 1111 on R1, R2, and R4 are excluded from the universal reachability requirement.
- In this exercise, R8 is used for backbone router simulation. Delivery of IP packets from R8 to the destinations beyond 172.16.1.0/24 is not necessary. However, unless explicitly specified otherwise, the routes advertised by R8 need to be present in the routing tables of other routers and traffic from other routers need to be able to reach R8.
- Unless explicitly specified otherwise, addresses and networks that are advertised in the "BGP" section need to be reachable by all BGP routers but do not need to be reachable by routers that use only IGP.
- Do not create new interfaces to fulfill IGP requirements, and do not summarize unless you are explicitly asked to do so.
- Do not introduce any new IPv4 addresses.
- Use conventional routing algorithms only.
- 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 172.16.0.0/16.

All IP addresses in this lab belong to the 172.16.0.0/16 address space, except for prefixes that are used in the "BGP Section" section.

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.

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 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 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 other 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. DMVPN Communications

Configure the mGRE Tunnel124 interfaces on R1, R2, and R4 according to the scenario requirements:

R1:
```
!
interface Tunnel124
 ip address 172.16.124.1 255.255.255.0
 tunnel source Loopback1
 tunnel mode gre multipoint
 tunnel key 10
!
```

R2:
```
!
interface Tunnel124
 ip address 172.16.124.2 255.255.255.0
 tunnel source Loopback1
 tunnel mode gre multipoint
 tunnel key 10
!
!
```

R4:
```
!
interface Tunnel124
 ip address 172.16.124.4 255.255.255.0
 tunnel source Loopback1
 tunnel mode gre multipoint
 tunnel key 10
!
```

Note that the **tunnel key 10** command is used to configure the tunnel key in this answer key. Because the lab does not specify the tunnel key value, you can use any number but it must match between the tunnel endpoints.

Configure the NHRP and DMVPN on R1, R2, and R4 according to the scenario requirements:

R1:
```
!
interface Tunnel124
 ip address 172.16.124.1 255.255.255.0
 no ip redirects
 ip nhrp network-id 10
 tunnel source Loopback1
 tunnel mode gre multipoint
 tunnel key 10
!
```

R2:
```
interface Tunnel124
 ip address 172.16.124.2 255.255.255.0
 no ip redirects
 ip nhrp map 172.16.124.1 1.1.1.1
 ip nhrp network-id 10
 ip nhrp nhs 172.16.124.1
 tunnel source Loopback1
 tunnel mode gre multipoint
 tunnel key 10
!
```

R4:

```
!
interface Tunnel124
 ip address 172.16.124.4 255.255.255.0
 no ip redirects
 ip nhrp map 172.16.124.1 1.1.1.1
 ip nhrp network-id 10
 ip nhrp nhs 172.16.124.1
 tunnel source Loopback1
 tunnel mode gre multipoint
 tunnel key 10
!
!
```

Note that R1 is defined as an NHS and the NHRP mapping for NHS is done on the NHRP spokes R2 and R4. Also, the DMVPN network ID is defined on all DMVPN routers with the **ip nhrp network-id 10** command.

Verify NHRP registrations on the NHS R1:

```
R1#show ip nhrp
172.16.124.2/32 via 172.16.124.2
   Tunnel124 created 00:02:14, expire 01:57:45
   Type: dynamic, Flags: unique registered
   NBMA address: 1.1.1.2
172.16.124.4/32 via 172.16.124.4
   Tunnel124 created 00:02:05, expire 01:57:54
   Type: dynamic, Flags: unique registered
   NBMA address: 1.1.1.4
R1#
```

Verify the DMVPN connectivity. Here is an example on R2:

```
R2#ping 172.16.124.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.124.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 8/8/9 ms
R2#ping 172.16.124.4
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.124.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 16/17/18 ms
R2#
```

Note that the spoke R2 can ping the hub R1 and the other spoke R4.

Note	The Mentor Guide engine in the web portal can help you use Cisco IOS Software commands to see a comprehensive view of the configuration tasks in this section engine. With the Mentor Guide engine, you can enter more than Cisco IOS Software commands as well a collection of proprietary commands such as **show all**.

2. 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,” “Switch-to-Router Connections,” and “Switch-to-Switch Connections” tables to analyze the VLAN’s propagation in this lab.

See the following diagram for the VLAN layout.

VLAN Propagation Diagram

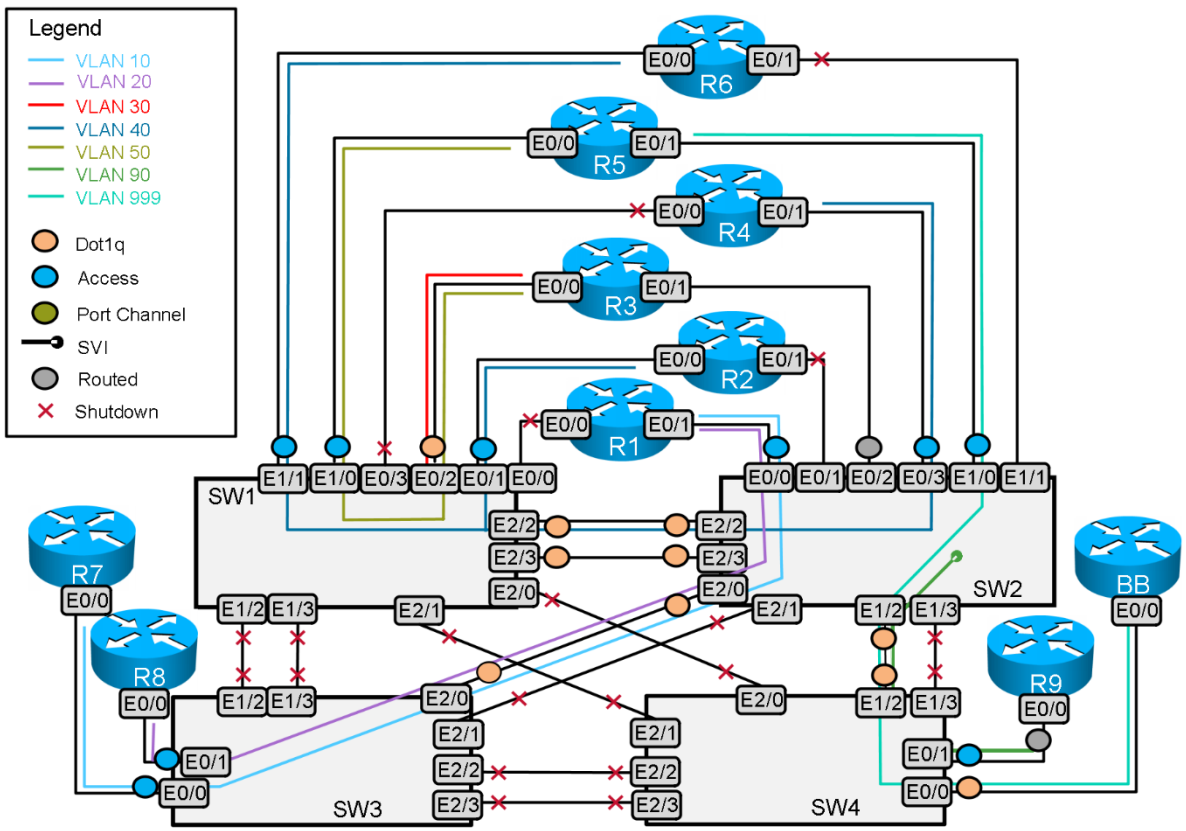

Carefully review the entire scenario. Closely examine the supplied diagram and any associated tables. Determine how you need to configure VTP, how to configure ports that are assigned as trunks, and how to configure ports that are assigned as simply static VLAN ports. Use the **switchport mode access**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

Issue: Configure rapid convergence based on the IEEE 802.1w standard with minimal extra configuration.
Each VLAN on the trunks between the switches should run its own spanning-tree instance.

Solution:

This issue can be solved by configuring rapid Per VLAN Spanning Tree Plus (PVST+). The rapid-PVST+ uses substantially the same configuration as PVST+, and the switch needs only minimal extra configuration. The benefit of rapid PVST+ is that a large PVST+ install base can be migrated to rapid PVST+ without the need to learn the complexities of the Multiple Spanning Tree Protocol (MSTP) configuration and without the need to reprovision the network. In rapid-PVST+ mode, each VLAN runs its own spanning-tree instance up to the maximum supported.

```
spanning-tree mode rapid-pvst
```

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 Cisco IOS Software commands as well a collection of proprietary commands such as **show all**.

3. IP Addresses Configuration

Issue: Configure SW2 to supply the IP addresses on subnet 172.16.90.0/24 to devices on VLAN 90. Assign only IP addresses 172.16.90.3, 172.16.90.5, and 172.16.90.7. Devices on VLAN 90 should use 172.16.90.1 as the default router.

Solution:

This DHCP configuration requirement is relatively simple.

SW2

```
ip dhcp excluded-address 172.16.90.1 172.16.90.2
ip dhcp excluded-address 172.16.90.4
ip dhcp excluded-address 172.16.90.6
ip dhcp excluded-address 172.16.90.8 172.16.90.254
!
ip dhcp pool VLAN90
   network 172.16.90.0 255.255.255.0
   default-router 172.16.90.1
!
```

SW2

```
interface Vlan90
 ip address 172.16.90.1 255.255.255.0
!
```

Note Note the interaction between this requirement and later requirements in the “Security” section.

Issue: Assign IP addresses from the 172.16.1.0/24 subnet to R1, R7, and the R8 switch virtual interface (SVI).

Solution:

The core issue involved with this configuration requirement is how to make devices attached to two separate VLANs appear on the same IP subnet. R1 and R7 are assigned to VLAN 10 while R1 and R8 are assigned to VLAN 20. Fulfill this requirement by configuring integrated routing and bridging (IRB) on router R1. Assign both of the Ethernet subinterfaces on R1 to a bridge

group. Then enable IRB and create a bridge-group virtual interface (BVI) on R1 using the 172.16.1.0/24 subnet. Do not assign an IP address to the Ethernet subinterfaces. Do not forget to enter the global configuration command **bridge 1 route ip** on R1 to allow bridged IP packets to get forwarded to the BVI.

R1

```
bridge irb

bridge 1 protocol ieee
bridge 1 route ip

interface Ethernet0/1
 no ip address
!
interface Ethernet0/1.10
 encapsulation dot1Q 10
 bridge-group 1
!
interface Ethernet0/1.20
 encapsulation dot1Q 20
 bridge-group 1
!
!
interface BVI1
 ip address 172.16.1.1 255.255.255.0
!
```

Verify that the bridge is active on R1 and is forwarding by using the **show bridge** command:

```
R1#show bridge

Total of 300 station blocks, 298 free
Codes: P - permanent, S - self

Bridge Group 1:

    Address       Action   Interface       Age   RX count   TX count
aabb.cc00.0b00   forward   Et0/1.10          2          6          4
aabb.cc00.0c00   forward   Et0/1.20          2          5          4
R1#
```

Use the **show spanning-tree** command to display additional parameters of spanning tree:

```
R1#show spanning-tree

 Bridge group 1 is executing the ieee compatible Spanning Tree protocol
  Bridge Identifier has priority 32768, address aabb.cc00.0110
  Configured hello time 2, max age 20, forward delay 15
  We are the root of the spanning tree
  Topology change flag not set, detected flag not set
  Number of topology changes 3 last change occurred 00:05:31 ago
          from Ethernet0/1.20
  Times:  hold 1, topology change 35, notification 2
          hello 2, max age 20, forward delay 15
  Timers: hello 0, topology change 0, notification 0, aging 300

 Port 17 (Ethernet0/1.10) of Bridge group 1 is forwarding
   Port path cost 100, Port priority 128, Port Identifier 128.17.
   Designated root has priority 32768, address aabb.cc00.0110
   Designated bridge has priority 32768, address aabb.cc00.0110
   Designated port id is 128.17, designated path cost 0
   Timers: message age 0, forward delay 0, hold 0
   Number of transitions to forwarding state: 1
   BPDU: sent 323, received 1
```

```
 Port 18 (Ethernet0/1.20) of Bridge group 1 is forwarding
   Port path cost 100, Port priority 128, Port Identifier 128.18.
   Designated root has priority 32768, address aabb.cc00.0110
   Designated bridge has priority 32768, address aabb.cc00.0110
   Designated port id is 128.18, designated path cost 0
   Timers: message age 0, forward delay 0, hold 0
   Number of transitions to forwarding state: 1
   BPDU: sent 341, received 22

R1#
```

Verify connectivity between R1, R7, and R8:

```
R1#ping 172.16.1.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.1.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 4/4/5 ms
R1#ping 172.16.1.10
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.1.10, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 2/4/5 ms
R1#

R7#ping 172.16.1.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.1.2, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
R7#
R8#ping 172.16.1.10
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.1.10, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
R8#
```

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 Cisco IOS Software commands as well a collection of proprietary commands such as **show all**.

4. IPv4 OSPF

Issue: On R3 and R5, assign loopback interfaces with the IP address 172.16.60.1/28 and 172.16.105.0/24 to OSPF area 44.

Solution:

R3 does not have a direct connection to area 0. It has a connection to area 33. If you assign an interface on R3 to an area other than 33, you will need to configure a virtual link with area 33 as the transit area. Then you will extend OSPF area 0 to R5 by configuring another virtual link through transit area 44.

Issue: Use the OSPF network type nonbroadcast for the 172.16.124.0/24 and 172.16.13.0 subnets.

Solution:

Remember that the 172.16.124.0/24 subnet is configured on a DMVPN hub-and-spoke topology. Because all OSPF packets have a Time to Live (TTL) = 1, OSPF spoke routers will never communicate with other spoke routers. Therefore, no spoke routers can become either a designated router (DR) or backup designated router (BDR). To ensure that this situation never happens, set the OSPF priority to 0 on the spoke routers R2 and R4 at the DMVPN interface level. At hub router R1, enter in two neighbor statements, one for R2 and the second for R4.

Issue: On R3, have the VLAN 30 subnet advertised via OSPF without including the subnet as one of your OSPF networks. Make sure that it is viewed as a type 1 route by OSPF and that the network can be reached from everywhere.

Solution:

Configure redistribution of the connected subnets into OSPF and use the **route-map** command to set the metric.

Configure OSPF on R1, R2, R3, R4, and R5:

R1:

```
!
interface Loopback101
 ip address 172.16.101.1 255.255.255.0
 ip ospf network point-to-point
!
interface Tunnel124
 ip address 172.16.124.1 255.255.255.0
 no ip redirects
 ip nhrp network-id 10
 ip ospf network non-broadcast
 tunnel source Loopback1
 tunnel mode gre multipoint
 tunnel key 10
!
interface Serial1/0
 ip address 172.16.13.1 255.255.255.0
 ip ospf network non-broadcast
 ip ospf priority 0
!
!
router ospf 100
 router-id 172.16.101.1
 area 33 virtual-link 172.16.103.1
 network 172.16.13.0 0.0.0.255 area 33
 network 172.16.101.0 0.0.0.255 area 0
 network 172.16.124.0 0.0.0.255 area 0
 neighbor 172.16.124.4
 neighbor 172.16.124.2
!
```

R2:

```
interface Tunnel124
 ip address 172.16.124.2 255.255.255.0
 no ip redirects
 ip nhrp map 172.16.124.1 1.1.1.1
 ip nhrp network-id 10
 ip nhrp nhs 172.16.124.1
```

```
 ip ospf network non-broadcast
 ip ospf priority 0
 tunnel source Loopback1
 tunnel mode gre multipoint
 tunnel key 10

!
router ospf 100
 network 172.16.124.0 0.0.0.255 area 0
!
```

R3:

```
interface Loopback60
 ip address 172.16.60.1 255.255.255.240
 ip ospf network point-to-point
!
!
interface Serial1/0
 ip address 172.16.13.3 255.255.255.0
 ip ospf network non-broadcast
!

router ospf 100
 router-id 172.16.103.1
 area 33 virtual-link 172.16.101.1
 area 44 virtual-link 172.16.105.1
 redistribute connected subnets route-map CON2OSP
 network 172.16.13.0 0.0.0.255 area 33
 network 172.16.35.0 0.0.0.255 area 44
 network 172.16.60.0 0.0.0.15 area 44
 neighbor 172.16.13.1
!
ip prefix-list P50 seq 5 permit 172.16.50.0/24
!
route-map CON2OSP permit 10
 match ip address prefix-list P50
 set metric-type type-1
!
route-map CON2OSP permit 20
!
```

R4:

```
interface Tunnel124
 ip address 172.16.124.4 255.255.255.0
 no ip redirects
 ip nhrp map 172.16.124.1 1.1.1.1
 ip nhrp network-id 10
 ip nhrp nhs 172.16.124.1
 ip ospf network non-broadcast
 ip ospf priority 0
 tunnel source Loopback1
 tunnel mode gre multipoint
 tunnel key 10
!
!
router ospf 100
 network 172.16.124.0 0.0.0.255 area 0
!
```

R5:

```
interface Loopback105
 ip address 172.16.105.1 255.255.255.0
 ip ospf network point-to-point
!
```



```
!
router ospf 100
 router-id 172.16.105.1
 area 44 virtual-link 172.16.103.1
 network 172.16.35.0 0.0.0.255 area 44
 network 172.16.105.0 0.0.0.255 area 44
!
```

Use the **show ip route** command to verify OSPF routing table. Here is an example on R2:

```
R2#show ip route ospf
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       a - application route
       + - replicated route, % - next hop override

Gateway of last resort is not set

      1.0.0.0/8 is variably subnetted, 6 subnets, 2 masks
O        1.1.1.1/32 [110/65] via 1.1.12.1, 00:38:38, Serial1/0
O        1.1.1.4/32 [110/129] via 1.1.12.1, 00:38:38, Serial1/0
O        1.1.14.0/24 [110/128] via 1.1.12.1, 00:38:38, Serial1/0
      172.16.0.0/24 is subnetted, 1 subnets
O E2     172.16.103.0 [110/20] via 172.16.124.1, 00:18:07, Tunnel124
      172.16.0.0/16 is variably subnetted, 13 subnets, 3 masks
O IA     172.16.13.0/24 [110/1064] via 172.16.124.1, 00:23:36, Tunnel124
O E2     172.16.31.0/24 [110/20] via 172.16.124.1, 00:18:07, Tunnel124
O IA     172.16.35.0/24 [110/1074] via 172.16.124.1, 00:23:21, Tunnel124
O E1     172.16.50.0/24 [110/1084] via 172.16.124.1, 00:18:37, Tunnel124
O IA     172.16.60.0/28 [110/1065] via 172.16.124.1, 00:23:21, Tunnel124
O        172.16.101.0/24 [110/1001] via 172.16.124.1, 00:36:48, Tunnel124
O IA     172.16.105.0/24 [110/1075] via 172.16.124.1, 00:00:59, Tunnel124
R2#
```

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 Cisco IOS Software commands as well a collection of proprietary commands such as **show all**.

5. IPv4 RIP

Issue: Do not broadcast or multicast RIP updates on VLAN 40 for security reasons.

Solution:

If you are instructed not to use broadcast or multicast RIP updates, configure RIP to unicast its updates. Make the interface passive and then configure neighbor statements for every device that needs to receive the RIP updates. In this scenario, you should configure two neighbor statements on each VLAN 40 RIP-speaking router.

R2 (example):

```
router rip
 version 2
 passive-interface default
 network 172.16.0.0
```

```
 neighbor 172.16.26.6
 neighbor 172.16.26.4
!
```

Enable debugging for IP packets on R2 (for example); you will see that RIP sends unicast updates on VLAN 40:

```
IP: s=172.16.26.2 (local), d=172.16.26.6 (Ethernet0/0), len 372, sending
    UDP src=520, dst=520
IP: s=172.16.26.2 (local), d=172.16.26.4 (Ethernet0/0), len 372, sending
    UDP src=520, dst=520
```

Issue: Routers R2 and R4 should prefer the RIP path when transmitting traffic between their respective loopback interfaces

Solution:

This issue will be discussed in the "IPv4 Route Redistribution" section.

Configure RIPv2 on R2, R4, and R6:

R2:

```
router rip
 version 2
 passive-interface default
 network 172.16.0.0
 neighbor 172.16.26.4
 neighbor 172.16.26.6
!
```

R4:

```
router rip
 version 2
 passive-interface default
 network 172.16.0.0
 neighbor 172.16.26.2
 neighbor 172.16.26.6
!
```

R6:

```
router rip
 version 2
 passive-interface default
 network 172.16.0.0
 neighbor 172.16.26.4
 neighbor 172.16.26.2
!
```

Use the **show ip route** command to verify the RIP routing table on R6:

```
R6#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 not set

      172.16.0.0/16 is variably subnetted, 9 subnets, 2 masks
R        172.16.102.0/24 [120/1] via 172.16.26.2, 00:00:13, Ethernet0/0
R        172.16.104.0/24 [120/1] via 172.16.26.4, 00:00:09, Ethernet0/0
R        172.16.124.0/24 [120/1] via 172.16.26.4, 00:00:09, Ethernet0/0
                         [120/1] via 172.16.26.2, 00:00:13, Ethernet0/0
R6#
```

Configure RIPv2 on R3 and SW2.

R3:

```
router rip
 version 2
 passive-interface default
 no passive-interface Ethernet0/1
 network 172.16.0.0
```

SW2:

```
router rip
 version 2
 passive-interface default
 no passive-interface Ethernet0/2
 network 172.16.0.0
!
```

Verify the RIP routing table on SW2:

```
SW2#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
       + - replicated route, % - next hop override

Gateway of last resort is not set

      172.16.0.0/16 is variably subnetted, 10 subnets, 3 masks
R        172.16.13.0/24 [120/1] via 172.16.31.3, 00:00:16, Ethernet0/2
R        172.16.35.0/24 [120/1] via 172.16.31.3, 00:00:16, Ethernet0/2
R        172.16.50.0/24 [120/1] via 172.16.31.3, 00:00:16, Ethernet0/2
R        172.16.60.0/28 [120/1] via 172.16.31.3, 00:00:16, Ethernet0/2
SW2#
```

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 Cisco IOS Software commands as well a collection of proprietary commands such as **show all**.

6. IPv4 EIGRP

Issue: Allow networks 192.168.2.0, 192.168.3.0, and 172.16.107.0 to be accepted into R1 from the router R8.

Solution:

On R1, configure an access list allowing only the two listed prefixes to be accepted by Enhanced Interior Gateway Routing Protocol (EIGRP). Apply the access list to a **distribute-list in** command referencing the BVI interface:

```
router eigrp 20
 network 172.16.1.0 0.0.0.255
 distribute-list 10 in BVI1
!
access-list 10 permit 0.0.2.0 255.255.253.0
```

Look at how access control list (ACL) 10 was constructed. R1 is receiving the following networks:

EIGRP Filter

Subnet	Action	Third Octet (dec)	Third Octet (bin)
172.16.107.0/24	Permit	107	01101011
192.168.1.0/24	Deny	1	00000001
192.168.2.0/24	Permit	2	00000010
192.168.3.0/24	Permit	3	00000011
192.168.4.0/24	Deny	4	00000100
192.168.5.0/24	Deny	5	00000101
		Address third octet (bin):	00000010
		Mask third octet (bin):	11111101

The networks that we need to permit and deny can be distinguished by a single bit within the third octet. To construct the filter, we will set this bit only in the IP address part of the ACL and clear it in the mask. As a result the router will look at only that bit (setting the mask bit to **1** tells the router to ignore the value during a match) and will permit networks only where the bit is set. We also will set a mask on the first two octets to all 1, which will tell the router to ignore the content of the first two octets.

Note This filter does not perform a precise match of three networks only. If R8 starts advertising additional networks, R1 may accept them.

Issue the **show ip route eigrp 20** command on R1:

```
R1#show ip route eigrp 20
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

      172.16.0.0/16 is variably subnetted, 15 subnets, 3 masks
D        172.16.107.0/24 [90/512000] via 172.16.1.2, 00:05:33, BVI1
D EX  192.168.2.0/24 [170/409600] via 172.16.1.2, 00:05:33, BVI1
```



```
D EX  192.168.3.0/24 [170/409600] via 172.16.1.2, 00:05:33, BVI1
R1#
```

Issue: Do not have R1 send EIGRP updates.

Solution:

At first, this issue seems like a passive interface configuration requirement. However, a passive EIGRP interface will not transmit any EIGRP hellos. If no hello packets are generated, the EIGRP speaker will never form an adjacency with another EIGRP speaker. If no adjacency is formed, then the EIGRP speaker will not receive any routing updates. R1 must receive updates from R8. Therefore, configuring the **passive-interface** command is unacceptable.

The solution to this problem could have been to configure a distribute list that denies all routes and to apply the distribute list to the BVI interface. However, such configuration is explicitly prohibited in this scenario.

Another option to fulfill this configuration requirement is to configure the **eigrp stub receive-only** command under the EIGRP routing process on router R1. Use this command to configure R1 to silently listen to R8 without advertising any updates to R8.

Note that the preceding configuration will result in R8 not receiving any EIGRP routes from R1, which will impact reachability. The "Restrictions and Goals" section provides for relaxation of reachability requirements.

R1:

```
router eigrp 20
 eigrp stub receive-only
!
```

Issue: Configure EIGRP 10 between R1 and R7.

Solution:

Because the **stub receive-only** command is configured in the R1 EIGRP AS 20 routing process, R7 should receive the EIGRP updates in AS 10. The routing information from the EIGRP AS 20 and OSPF will be redistributed into EIGRP AS 10 in the redistribution section. An interesting feature of this configuration is that the 172.16.1.0/24 prefix assigned to BVI1 on R1 resides in both EIGRP AS 10 and EIGRP AS 20:

```
R1#show ip eigrp topology 172.16.1.0 255.255.255.0
EIGRP-IPv4 Topology Entry for AS(20)/ID(172.16.101.1) for 172.16.1.0/24
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 384000
  Descriptor Blocks:
  0.0.0.0 (BVI1), from Connected, Send flag is 0x0
      Composite metric is (384000/0), route is Internal
      Vector metric:
        Minimum bandwidth is 10000 Kbit
        Total delay is 5000 microseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 0
        Originating router is 172.16.101.1
EIGRP-IPv4 Topology Entry for AS(10)/ID(172.16.101.1) for 172.16.1.0/24
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 384000
  Descriptor Blocks:
  0.0.0.0 (BVI1), from Connected, Send flag is 0x0
      Composite metric is (384000/0), route is Internal
```

```
      Vector metric:
        Minimum bandwidth is 10000 Kbit
        Total delay is 5000 microseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 0
        Originating router is 172.16.101.1
R1#
```

Configure EIGRP on R1, R7, and R8:

R1:

```
router eigrp 20
 distribute-list 10 in BVI1
 network 172.16.1.0 0.0.0.255
 eigrp stub receive-only
!
!
router eigrp 10
 network 172.16.1.0 0.0.0.255
!
access-list 10 permit 0.0.2.0 255.255.253.0
```

R7:

```
router eigrp 10
 network 172.16.1.0 0.0.0.255
 network 172.16.110.0 0.0.0.255
!
```

R8:

```
router eigrp 20
 network 172.16.1.0 0.0.0.255
 network 172.16.107.0 0.0.0.255
 redistribute connected metric 10000 100 255 1 1500
!
```

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 Cisco IOS Software commands as well a collection of proprietary commands such as **show all**.

7. IPv4 Route Redistribution

Before you examine the specific issues related to configuring each of the IGPs involved in this scenario, you will survey the entire topology and determine how all the different IGPs will interoperate. Performing such a survey forces you to consider the issues related to route redistribution. When you evaluate a single internetwork topology that contains multiple routing protocols, a good starting point of analysis is to determine if more than one direct or indirect connecting point is between two routing protocols. If only one connecting point is between two routing protocols, providing connectivity between them is relatively simple. If there are two or more connecting points, then providing connectivity between the two routing protocols can be complicated. When two or more connecting points exist, you can use them to provide redundancy and for load balancing and optimum path selection. However, you must also at least ensure that no routing loops exist and, whenever possible, that no suboptimal paths are selected.

The redistribution points are R2, R4, R1, R3, and R8.

There is a two-point mutual redistribution between RIP and OSPF in this scenario. This particular topology represents an academic redistribution problem in which loops and prefix loss will occur unless special precautions are taken.

To ensure that RIP and OSPF redistribute flawlessly and seamlessly, RIP must not redistribute its own prefixes back from OSPF. Also, RIP should set the administrative distance for its own routes to 109 (or any value lower than the OSPF distance 110).

This process will ensure that RIP will not lose its own prefixes and that the redistributing router will not install OSPF external prefixes instead of RIP native prefixes in its local routing table.

A loop avoidance technique also is applied in this scenario. If you shut down loopback 30.1 on router R6 and assume that R4 is the Autonomous System Boundary Router (ASBR), R1 will point the route to the Loopback30 network to R4, R4 will point to R2, and R2 will point back to R1. A routing loop is formed. This potential routing loop condition needs to be prevented. RIP must not accept its native prefixes from OSPF and send them back to OSPF on the ASBR. RIP should be configured to not redistribute its native prefixes from OSPF.

R1 redistributes all routes between OSPF and EIGRP.

R3 performs mutual redistribution between RIP and OSPF. One of the connected networks is redistributed as E1, others as E2. Therefore, a route map will have two entries. The first entry will be **set metric-type type-1** (matching the router we need to advertise with this metric type) and the second entry will permit all other routes. The **set metric** statement is not required for the second route map entry, because type-2 is the default.

The following table provides a useful summary of which prefixes were imported into a given routing protocol. An empty permit column for a given routing protocol indicates that no prefixes were redistributed into the routing protocol. This result represents that the routing protocol is involved in one-way redistribution.

IPv4 IGP Redistribution

Redist point	Into RIP		Into OSPF		Into EIGRP 10		Into EIGRP 20	
	PERMIT	**DENY**	**PERMIT**	**DENY**	**PERMIT**	**DENY**	**PERMIT**	**DENY**
R1			All EIGRP 10 and 20 routes		All EIGRP 10 and 20 routes; all OSPF routes		None	
R2	All OSPF routes	RIP native routes	RIP native routes					
R3	All OSPF routes		All RIP routes					
R4	All OSPF routes	RIP native routes	RIP native routes					
R8							Connected networks	

Configure the IGP route redistribution on R1, R2, R3, and R4. The redistribution of the connected networks on R8 is already configured in the EIGRP section.

R1:

```
!
router eigrp 10
redistribute eigrp 20
redistribute ospf 100 metric 10000 100 255 1 1500
!
!
router ospf 100
 redistribute eigrp 20 subnets
 redistribute eigrp 10 subnets
!
```

R2:

```
!
router ospf 100
 redistribute rip metric 1 subnets route-map RIP-->OSPF
!
router rip
 redistribute ospf 100 metric 1 route-map OSPF-->RIP
 distance 109 0.0.0.0 255.255.255.255 RIP-distance
!
!
ip access-list standard RIP-->OSPF
 permit 172.16.30.0
 permit 172.16.26.0
 permit 172.16.104.0
 permit 172.16.106.0
 permit 172.16.102.0
ip access-list standard RIP-distance
 permit 172.16.30.0
 permit 172.16.104.0
 permit 172.16.106.0
!
!
route-map RIP-->OSPF permit 10
 match ip address RIP-->OSPF
!
route-map OSPF-->RIP deny 10
 match ip address RIP-->OSPF
!
route-map OSPF-->RIP permit 20
!
!
```

R3:

```
!
router ospf 100
 redistribute connected subnets route-map CON2OSP
 redistribute rip subnets
!
router rip
redistribute ospf 100 metric 1
!
!
ip prefix-list P50 seq 5 permit 172.16.50.0/24
!
```



```
route-map CON2OSP permit 10
 match ip address prefix-list P50
 set metric-type type-1
!
route-map CON2OSP permit 20
!
```

R4:

```
router ospf 100
 redistribute rip metric 1 subnets route-map RIP-->OSPF
!
router rip
 redistribute ospf 100 metric 1 route-map OSPF-->RIP
 distance 109 0.0.0.0 255.255.255.255 RIP-distance
!
!
ip access-list standard RIP-->OSPF
 permit 172.16.30.0
 permit 172.16.26.0
 permit 172.16.104.0
 permit 172.16.106.0
 permit 172.16.102.0
ip access-list standard RIP-distance
 permit 172.16.30.0
 permit 172.16.106.0
 permit 172.16.102.0
!
!
route-map RIP-->OSPF permit 10
 match ip address RIP-->OSPF
!
route-map OSPF-->RIP deny 10
 match ip address RIP-->OSPF
!
route-map OSPF-->RIP permit 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 Cisco IOS Software commands as well a collection of proprietary commands such as **show all**.

8. Border Gateway Protocol

BGP Diagram

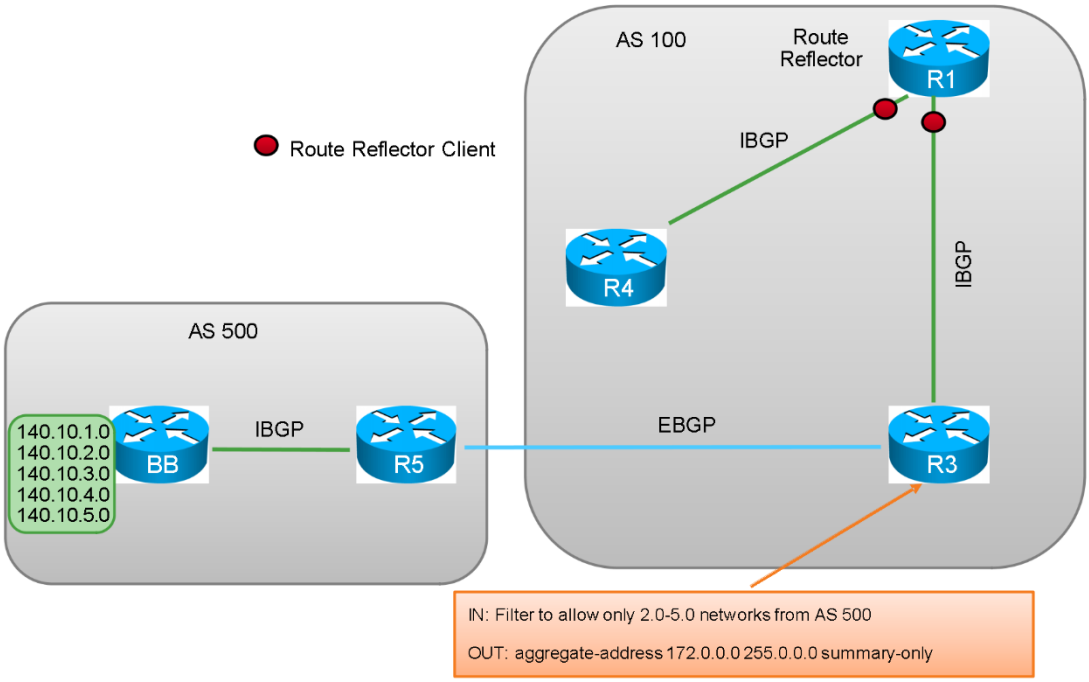

Issue: Configure R3 to accept only the following networks from AS 500:

- 140.10.2.0/24
- 140.10.3.0/24
- 140.10.4.0/24
- 140.10.5.0/24

Use the ACL with no more than two lines. Filtering should remain accurate even if additional networks are advertised by AS 500.

Solution:

You will start by considering a requirement that the filter should remain accurate even if additional networks are advertised. Unlike in EIGRP filtering described earlier, you have to construct a filter that would permit these four networks only.

You are also specifically instructed to use ACL and not a prefix list. Although filtering with extended ACLs is more flexible than with prefix lists, use of extended ACLs is much less intuitive (there are also performance advantages of prefix list use).

Extended ACLs when used for route filtering consist of two parts. The first part (what would be the IP source if you were to use ACL for security) defines the prefix. The second part (what would be the IP destination if you were to use ACL for security) defines the prefix mask.

Because all the prefixes have masks of /24, the second half of the ACL will read “host 255.255.255.0.”

Because you are permitted to use two lines, you can use "140.10.2.0 0.0.1.0" to match 140.10.2.0 and 140.10.3.0 (recall that setting a mask bit to 1 tells the Cisco IOS Software not to care about the content of that bit) and "140.10.4.0 0.0.1.0" to match 140.10.4.0 and 140.10.5.0.

The final ACL looks as follows:

```
access-list 100 permit ip 140.10.2.0 0.0.1.0 host 255.255.255.0
access-list 100 permit ip 140.10.4.0 0.0.1.0 host 255.255.255.0
```

In this example, **distribute-list in** is used as a variant and access list 100 controls which prefixes are allowed in:

```
router bgp 100
 address-family ipv4
 neighbor 172.16.35.5 distribute-list 100 in
!
```

Issue **show ip bgp** command on R3:

```
R3#show ip bgp
   Network          Next Hop            Metric LocPrf Weight Path
*> 140.10.2.0/24    172.16.35.5              0             0 500 i
*> 140.10.3.0/24    172.16.35.5              0             0 500 i
*> 140.10.4.0/24    172.16.35.5              0             0 500 i
*> 140.10.5.0/24    172.16.35.5              0             0 500 i
*> 172.0.0.0/8      0.0.0.0                               32768 i
s> 172.16.50.0/24   0.0.0.0                  0            32768 i
```

Issue: The networks learned by R3 from AS 500 should be present in the BGP table on R4.

Solution:

R4 is the third IBGP speaker within AS 100. In order for R4 to see the 140.10.*.* prefixes in its BGP table and to fulfill the partial mesh requirement, configure a route reflector within AS 100. Without a route reflector or confederation, IBGP speakers need to be fully meshed.

R1

```
router bgp 100
 neighbor 172.16.13.3 route-reflector-client
 neighbor 172.16.124.4 route-reflector-client
!
```

Issue: R4 should have only a single BGP route 140.10.2.0/24 in its IPv4 routing table; no other BGP routes are allowed on R4. Do not introduce modifications in the R4 BGP table to meet this requirement.

Solution:

Because we are required not to change the BGP table, we need to find a solution of filtering routing information between BGP and IP tables. Unlike in other protocols (such as OSPF), in BGP a filter cannot be attached in between the tables. BGP does, however, support the **distance** command, although this command usually is used to change relative preference between protocols. Setting the administrative distance to 255 prevents a router from being installed into the IP routing table.

Issue the **show ip bgp** and **show ip route bgp** commands on R4. Only 140.10.2.0 will be installed; other prefixes are filtered by the **distance** command and access list:

R4:

```
router bgp 100
 address-family ipv4
```

```
 distance 255 0.0.0.0 255.255.255.255 BGP-distance
!
ip access-list standard BGP-distance
 deny   140.10.2.0
 permit any

R4#show ip bgp
BGP table version is 34, local router ID is 172.16.104.1
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
*>i140.10.2.0/24    172.16.35.5              0    100      0 500 i
r>i140.10.3.0/24    172.16.35.5              0    100      0 500 i
r>i140.10.4.0/24    172.16.35.5              0    100      0 500 i
r>i140.10.5.0/24    172.16.35.5              0    100      0 500 i

R4#show ip route bgp
     140.10.0.0/24 is subnetted, 1 subnets
B       140.10.2.0 [200/0] via 172.16.35.5, 00:02:55
```

Issue: Send an aggregate for 172.0.0.0/8 to AS 500 and suppress all other routes.

Solution:

Configure an aggregate for 172.0.0.0/8 on R3 with the summary-only option. R3 must have a 172.x.x.x entry in its BGP table to be able to advertise the aggregate.

R3:

```
router bgp 100
 address-family ipv4
 network 172.16.50.0 mask 255.255.255.0
 aggregate-address 172.0.0.0 255.0.0.0 summary-only
!
```

Issue the **show ip bgp** command on R5:

```
R5#show ip bgp
   Network          Next Hop            Metric LocPrf Weight Path
*> 140.10.1.0/24    0.0.0.0                  0         32768 i
*> 140.10.2.0/32    0.0.0.0                  0         32768 i
*> 140.10.2.0/24    0.0.0.0                            32768 i
*> 140.10.3.0/24    0.0.0.0                  0         32768 i
*> 140.10.4.0/24    0.0.0.0                  0         32768 i
*> 140.10.5.0/24    0.0.0.0                  0         32768 i
*> 140.10.6.0/24    0.0.0.0                  0         32768 i
*> 172.0.0.0/8      172.16.35.3              0             0 100 i
```

The BGP configuration on R5, R3, R1, and R4 is as follows:

R5:

```
router bgp 500
 bgp log-neighbor-changes
 neighbor 10.1.1.1 remote-as 500
 neighbor 172.16.35.3 remote-as 100
!
```

R3:

```
router bgp 100
 bgp log-neighbor-changes
 network 172.16.50.0 mask 255.255.255.0
 aggregate-address 172.0.0.0 255.0.0.0 summary-only
 neighbor 172.16.13.1 remote-as 100
 neighbor 172.16.35.5 remote-as 500
```



```
 neighbor 172.16.35.5 distribute-list 100 in
!
!
access-list 100 permit ip 140.10.2.0 0.0.1.0 host 255.255.255.0
access-list 100 permit ip 140.10.4.0 0.0.1.0 host 255.255.255.0
!
```

R1:

```
router bgp 100
 bgp log-neighbor-changes
 neighbor 172.16.13.3 remote-as 100
 neighbor 172.16.13.3 route-reflector-client
 neighbor 172.16.124.4 remote-as 100
 neighbor 172.16.124.4 route-reflector-client
!
```

R4:

```
router bgp 500
router bgp 100
 bgp log-neighbor-changes
 neighbor 172.16.124.1 remote-as 100
 distance 255 0.0.0.0 255.255.255.255 BGP-distance
!
!
ip access-list standard BGP-distance
 deny   140.10.2.0
 permit any
```

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 Cisco IOS Software commands as well a collection of proprietary commands such as **show all**.

9. Router Maintenance

Issue: Configure R4 to supply configuration information to a new router that will be connected to VLAN 40. The new router should receive its configuration from the TFTP server 172.16.50.100 located on the VLAN 30. The new router will have IP address 172.16.26.100/24 and MAC address 0010.7be8.131d, and will use interface E0/1. The new router should send the request for the configuration R100.cfg via R2.

Solution:

Read the entire “Router Maintenance” section to learn how to use all the supplied information. The Autoinstall over Ethernet feature allows you to automate the router configuration. R4 should be configured as a DHCP server supplying all necessary information to the Autoinstall client (a new router). The client is on the same subnet with the DHCP server.

You can specify the address of the TFTP server by configuring the **option 150 ip** *X.X.X.X* command, where *X.X.X.X* is a given TFTP server IP address, in ip dhcp pool configuration mode:,.

R4:

```
no ip dhcp conflict logging
!
ip dhcp pool NEW-ROUTER
   option 150 ip 172.16.50.100
```

Client software determines the method a DCHP client uses to identify himself to a DHCP server and then the configuration required on the DHCP server to create a lease for such a client. This task specifies that the client is Cisco IOS Software Release 12.4. The client uses the client-

identifier field of DHCP request and constructs this field by concatenating "*null*cisco-", the MAC address, and the interface identifier, as illustrated in the following:

ASCII client ID	null	c	i	s	c	o	-	0	0	1	0	.	7	b	e	8	.	1	3	1	d	-	F	a	0	/	1
Hex conversion	00	63	69	73	63	6F	2D	30	30	31	30	2E	37	62	65	38	2E	31	33	31	64	2D	46	61	30	2F	31
Hex client ID	0063.6973.636F.2D30.3031.302E.3762.6538.2E31.3331.642D.4661.302F.31																										

Obtain the information on the constructed client ID from the router rather than by performing a manual conversion.

The best way to obtain the information is from the output of the **show**/**debug** messages on the DHCP server side, But the information can also be obtained on the client side. In the example that follows, an unused interface on R2 is used to obtain the client ID for this configuration requirement.

You start by selecting the same interface (E0/1) and assigning the same MAC address according to the requirements of the task If you want to test the DHCP communications between R2 and R4, do not forget to add the VLAN 30 configuration to ensure that E0/1 on R2 is connected to R4):

R2:

```
interface Ethernet0/1
 mac-address 0010.7be8.131d
 ip address dhcp
!
```

Turn on debugging of DHCP client activity and observe the output:

```
R2#deb dhcp detail
DHCP client activity debugging is on (detailed)
R2#
*Apr 25 01:25:17.664: DHCP: QScan: Timed out Selecting state
R2#%Unknown DHCP problem.. No allocation possible
*Apr 25 01:25:27.138: DHCP: Waiting for 5 seconds on interface Ethernet0/1
R2#
*Apr 25 01:25:32.142: DHCP: Try 2 to acquire address for Ethernet0/1
*Apr 25 01:25:32.149: DHCP: allocate request
*Apr 25 01:25:32.149: DHCP: zapping entry in DHC_PURGING state for Et0/1
*Apr 25 01:25:32.149: DHCP: deleting entry F0CA2538 0.0.0.0 from list
*Apr 25 01:25:32.149: Temp IP addr: 0.0.0.0  for peer on Interface: Ethernet0/1
*Apr 25 01:25:32.149: Temp  sub net mask: 0.0.0.0
*Apr 25 01:25:32.149:    DHCP Lease server: 0.0.0.0, state: 11 Purging
*Apr 25 01:25:32.149:    DHCP transaction id: 17E7
*Apr 25 01:25:32.149:    Lease: 0 secs,  Renewal: 0 secs,  Rebind: 0 secs
*Apr 25 01:25:32.150:    Next timer fires after: 00:00:26
*Apr 25 01:25:32.150:    Retry count: 0   Client-ID: cisco-0010.7be8.131d-Et0/1
*Apr 25 01:25:32.150:    Client-ID hex dump: 636973636F2D303031302E376265382E
*Apr 25 01:25:32.150:                        313331642D4574302F31
*Apr 25 01:25:32.150:    Hostname: R2
*Apr 25 01:25:32.150: DHCP: new entry. add to queue, interface Ethernet0/1
*Apr 25 01:25:32.150: DHCP: SDiscover attempt # 1 for entry:
*Apr 25 01:25:32.150: Temp IP addr: 0.0.0.0  for peer on Interface: Ethernet0/1
*Apr 25 01:25:32.150: Temp  sub net mask: 0.0.0.0
*Apr 25 01:25:32.150:    DHCP Lease server: 0.0.0.0, state: 3 Selecting
*Apr 25 01:25:32.150:    DHCP transaction id: 17E8
```

When you use this method, the leading "null" will not be shown in the debug output on the client side. You will have to manually add it. The client ID field needs to be dot-separated in groups of four hexadecimal digits (groups start at the beginning of the string):

```
   Client-ID hex dump: 636973636F2D303031302E376265382E313331642D4574302F31
↓
```



```
   client-identifier
0063.6973.636F.2D30.3031.302E.3762.6538.2E31.3331.642D.4574.302F.31
```

Configure the binding between a specified IP address and a MAC address using the **client-identifier** command in ip dhcp pool configuration mode.

R4:

```
ip dhcp pool NEW-ROUTER
   host 172.16.26.100 255.255.255.0
   client-identifier
0063.6973.636f.2d30.3031.302e.3762.6538.2e31.3331.642d.4661.302f.31
   bootfile R100.cfg
   option 150 ip 172.16.50.100
   default-router 172.16.26.2
```

Issue: Configure the 7.7.7.3/24 address on the E0/1 interface of R3 without changing any preexisting IP addresses; do not advertise to any routing protocol. Ensure that the workstation on the 7.7.7.0/24 private address space can reach the rest of the network by using a portion of the address space of the R3 primary E0/1 subnet.

Solution:

This task is a Network Address Translation (NAT) configuration requirement without ever explicitly mentioning NAT. The 7.7.7.3 address assigned to the R3 E0/1 interface will be assigned as a secondary IP address. This 7.7.7.0 will be the NAT inside address. The primary IP address of the E0/1 interface will be the NAT outside address.

Configure R3 E0/1 as the NAT inside interface:

```
interface Ethernet0/1
 ip address 7.7.7.3 255.255.255.0 secondary
 ip address 172.16.31.3 255.255.255.0
 ip nat inside
!
```

Configure other interfaces as the NAT outside interface:

```
interface Serial1/0
 ip address 172.16.13.3 255.255.255.0
 ip nat outside
!
interface Ethernet0/0.30
 encapsulation dot1Q 30
 ip address 172.16.50.3 255.255.255.0
 ip nat outside
!
interface Ethernet0/0.50
 encapsulation dot1Q 50
 ip address 172.16.35.3 255.255.255.0
 ip nat outside
!
```

Configure NAT translation and a list of source addresses:

```
access-list 50 permit 7.7.7.0 0.0.0.255
ip nat inside source list 50 interface E0/1 overload
```

Because this is a dynamic NAT, you will not see translations until packets start flowing. You can still check that the configuration is accepted by using **show ip nat statistics** command:

```
R3#show ip nat statistics
Total active translations: 1 (0 static, 1 dynamic; 1 extended)
Peak translations: 2, occurred 00:00:55 ago
Outside interfaces:
  Ethernet0/0.30, Ethernet0/0.50, Serial1/0
Inside interfaces:
```

```
  Ethernet0/1
Hits: 20  Misses: 0
CEF Translated packets: 10, CEF Punted packets: 0
Expired translations: 1
Dynamic mappings:
-- Inside Source
[Id: 1] access-list 50 interface Ethernet0/1 refcount 1

Total doors: 0
Appl doors: 0
Normal doors: 0
Queued Packets: 0
R3#
```

Use the **ping** command to verify connectivity via the NAT and the NAT translation table:

```
R3#ping 172.16.26.6 so 7.7.7.3
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.26.6, timeout is 2 seconds:
Packet sent with a source address of 7.7.7.3
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 15/16/17 ms
R3#
R3#
R3#ping 172.16.124.1 so 7.7.7.3
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.124.1, timeout is 2 seconds:
Packet sent with a source address of 7.7.7.3
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 8/8/9 ms

R3#show ip nat translations
Pro Inside global      Inside local       Outside local      Outside global
icmp 172.16.31.3:8     7.7.7.3:8          172.16.26.6:8      172.16.26.6:8
icmp 172.16.31.3:9     7.7.7.3:9          172.16.124.1:9     172.16.124.1:9
R3
```

Issue: Allow only Telnet to R7 and only from the R3 management loopback interface.

Solution:

Configuring R7 so that it can ping all interfaces in the scenario network is straightforward. The method to apply to allow Telnet access only from the R3 management loopback interface is not as clear. To accomplish this task, create an access list permitting only the R3 loopback interface. Then, apply the access list under the "line vty" mode with the **access-class 1 in** command:.

R7:

```
access-list 1 permit 172.16.103.1
!
line vty 0 4
 access-class 1 in
 privilege level 15
 transport input telnet
!
```

Use Telnet to R7 from R3 using Loopback103 as the source interface:

```
R3#telnet 172.16.1.10 /source-interface Loopback103
Trying 172.16.1.10 ... Open

------------------------------------------------
Cisco 360 R&S Workbook Labs
Product, POD location:  cierswbv5-ce-lab02-sc, SJ
```



```
Device:                 R7
-------------------------------------------------
R7#exit

[Connection to 172.16.1.10 closed by foreign host]
R3#

R3#telnet 172.16.10.1
Trying 172.16.10.1 ...
% Connection timed out; remote host not responding
R3#
```

Issue: Simulate traffic every 2 minutes by sending 100 test packets. Monitor the jitter of this simulated traffic from R5 to R1. Provide a minimal configuration on R1.
Provide two separate sets of simulated traffic from R5 to R1, with one set of traffic marked as precedence 4 and the second set of traffic marked as precedence 3.
For any type of IP service that needs to be configured for the task described, make sure that R1 is the server of any supporting service.

Solution:

Note	Cisco IOS Software syntax for IP Service Level Agreements (SLA) commands have changed multiple times, hence actual commands may be different on equipment you are using. See the Cisco IOS Software command reference for your release of software for correct syntax.

This is a Service Assurance Agent (SAA) configuration requirement. You can configure SAA or the Cisco IP SLA feature to simulate traffic on a periodic basis and measure delay variation between packets. Enter a collection of **rtr** or **ip sla** commands on the router that will be generating the traffic. In this specific configuration, these commands will be entered on R5. Only one global configuration command (**rtr responder** or **ip sla monitor responder**) needs to be configured on R1. This configuration fulfills the "minimal configuration on R1" requirement.

With the Cisco IOS IP SLA, you can specify type of service (ToS) settings for a configured set of test packets. Remember that you configure this feature with ToS values and not with IP precedence or differentiated services code point (DSCP) values. Therefore, you must be careful to map the ToS value to the correct IP precedence or DSCP value. The following table helps you perform this translation to meet the specific requirements of this scenario:

Type of Service (ToS) Octet

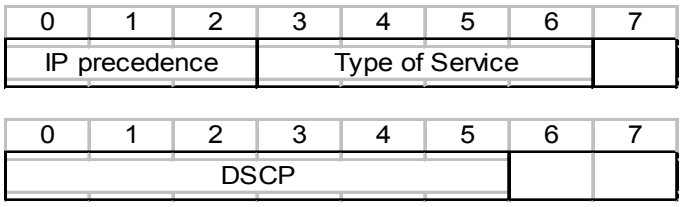

The table shows that you need to shift the IP precedence value 5 left to arrive to the ToS octet value. Accomplish the shift by multiplying the decimal value of the IP precedence by 32, as illustrated in the table:

ToS Octet and IP Precedence Decimal Values

Baseline ToS octet Decimal Value	128	64	32	16	8	4	2	1
IP Precedence Value	4	2	1					

The following table illustrates conversion between IP precedence and ToS octet values in binary and decimal presentations. (*Note that the conversion described here assumes that other bits of the ToS octet are set to zero.*)

IP Precedence ⇔ ToS Octet Conversion

IP precedence			ToS	
	Decimal	Binary	Binary	Decimal
Routine or Best Effort	0	000	00000000	0
Priority	1	001	00100000	32
Immediate	2	010	01000000	64
Flash	3	011	01100000	96
Flash Override	4	100	10000000	128
Critical	5	101	10100000	160
Internetwork Control	6	110	11000000	192
Network Control	7	111	11100000	224

The scenario directs you to mark a set of test packets with the IP precedence setting of 4. The table shows that an IP precedence value of 4 maps to a ToS value of 128.

The scenario also directs you to mark a separate set of test packets with the IP precedence setting of 3. The table shows that an IP precedence value of 3 maps to a ToS value of 96.

The following formulas summarize the conversion between ToS octet values and QoS bits:

```
ToS-Octet = IP-precedence * 32
ToS-Octet = DSCP * 4
```

R5:

```
ip sla 1
 udp-jitter 172.16.101.1 16387 num-packets 100
 tos 128
 frequency 120
ip sla schedule 1 life forever start-time now
ip sla 2
 udp-jitter 172.16.101.1 16388 num-packets 100
 tos 96
 frequency 120
ip sla schedule 2 life forever start-time now
```

R1

```
ip sla responder
```

Issue: Provide a mechanism that can be used on R1 to count the number of packets received on the S1/0 interface. Packet counts should be classified by the precedence setting of the received packets. Do not use any access lists to fulfill this requirement.

Solution:

One way to count the number of packets received on an interface on a per-IP precedence basis is by using the interface configuration command **ip accounting**.

R1:

```
interface Serial1/0
ip accounting precedence input
```

In a few moments verify the IP accounting on the S1/0 interface on R1:

```
R1#show interfaces s1/0 precedence
```



```
Serial1/0
  Input
    Precedence 3:  101 packets, 6484 bytes
    Precedence 4:  101 packets, 6484 bytes
    Precedence 6:  16 packets, 1326 bytes
R1#
```

Notice how only 101 packets are received. SAA and Cisco IOS IP SLA are configured to send exactly 100 packets. Consider the first packet as an SAA or Cisco IOS IP SLA control plane packet. This is how you end up with the amount of 101 packets. Do not forget to clear counters before experimenting.

Note Another way to mark IP SLA traffic with IP precedence is by using a QoS service policy.

This configuration task also directs your Network Time Protocol (NTP) configuration for SAA. Before SAA jitter values can be collected and processed, NTP must be running between the two test devices. The two SAA test devices, the generator and the responder, must be synchronized with NTP. Although the configuration requirements in this scenario are silent about this, you should know this if you have developed expertise in configuring the jitter testing feature of SAA. Now, that you know you need to configure NTP with the SAA configuration, you must determine whether to configure an NTP peer relationship or client/server relationship. Because the configuration task states that you must make R1 the "server for any supporting service," this directive applies to directly to your NTP configuration. Therefore, configure R1 as the NTP server and R5 as the NTP client.

R1:

```
ntp master
```

R5:

```
ntp server 172.16.101.1
```

Verify the NTP configuration on R5:

```
R5#show ntp  associations detail
172.16.101.1 configured, ipv4, our_master, sane, valid, stratum 8
ref ID 127.127.1.1    , time D5231252.A6A7F168 (18:24:18.651 PST Wed Apr 24
2013)
our mode client, peer mode server, our poll intvl 64, peer poll intvl 64
root delay 0.00 msec, root disp 3939.37, reach 1, sync dist 4134.01
delay 8.00 msec, offset 0.0000 msec, dispersion 189.44, jitter 0.97 msec
precision 2**10, version 4
assoc id 10487, assoc name 172.16.101.1
assoc in packets 6, assoc out packets 6, assoc error packets 0
org time 00000000.00000000 (16:00:00.000 PST Wed Dec 31 1899)
rec time D523125B.322D0EE0 (18:24:27.196 PST Wed Apr 24 2013)
xmt time D523125B.322D0EE0 (18:24:27.196 PST Wed Apr 24 2013)
filtdelay =     8.00    8.00    8.00    9.00    9.00    9.00    0.00    0.00
filtoffset =    0.00    0.00    0.00   -0.50   -0.50    0.50    0.00    0.00
filterror =     1.95    1.98    2.01    2.04    2.07    2.10 16000.0 16000.0
minpoll = 6, maxpoll = 10

R5#

R5#show clock detail
18:25:13.954 PST Wed Apr 24 2013
Time source is NTP
R5#
```

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 Cisco IOS Software commands as well a collection of proprietary commands such as **show all**.

10. Security

Issue: The administrator does not want any packet to be routed in the network based on the routing path carried in the IP packet.

Solution:

Disable source routing with the **no ip source-route** global configuration command. On some Cisco IOS Software, this command might be disabled by default.

Issue: Do not allow DHCP services**.**

Solution:

Disable DHCP services with the **no service dhcp** global configuration command.

Issue: Ensure that abnormally terminated TCP sessions are removed.

Solution:

By default, Cisco routers do not continually test whether a previously connected TCP endpoint is still reachable. If one end of a TCP connection idles out or terminates abnormally (such as through a reloads), the opposite end of the connection may still believe the session is available. These "orphaned" sessions use up valuable router resources. Attackers have been known to take advantage of this weakness to attack routers.

To remedy this situation, configure the routers to send periodic keepalive messages to ensure that the remote end of a session is still available. If the remote device fails to respond to the keepalive message, the sending router will clear the connection. This practice immediately frees router resources for other more important tasks. Keepalives are important because they help guard against orphaned sessions.

Apply the **service tcp-keepalives-in** and **service tcp-keepalives-out** commands in global configuration mode to enable TCP keepalives.

Issue: R4 should ignore and not reply to Bootstrap Protocol request packets received.

Solution:

To allow the DHCP server to selectively ignore and not reply to received Bootstrap Protocol (BOOTP) request packets, use the **ip dhcp bootp ignore** command.

Configuration on R5:

```
R5#show running-config | inc source-route|dhcp|tcp-keepalive
service tcp-keepalives-in
service tcp-keepalives-out
no service dhcp
no ip source-route
R5#
!
```


Issue: SW2 will provide security separation between the test segment and the rest of your network. Provide a solution to allow traffic sourced on the network 172.16.90.0/24 only from the selected hosts (.1, .3, .5, .7) to get into your network. Apply the solution on SW2. Do not use filtering techniques based on Layer 2 filtering. Use a minimal number of statements for this task.

Solution:

Notice that all hosts have odd numbers and are in the range from 1 to 7. In the binary form, all hosts can be represented as follows:

- 00000001 - .1
- 00000011 - .3
- 00000101 - .5
- 00000111 - .7

Notice that the first five digits from the left must be 0 and that the last digit must be 1. You do not need to be concerned about the second and third digits from the right; they can vary. This logic can be represented by the following base and wildcard:
172.16.90.1 0.0.0.6

This fulfills the requirement of matching the desired range of addresses with the minimum number of statements.

Apply your outbound access list on the Layer3 interface Ethenet0/2 on SW2.

SW2:

```
interface Ethernet0/2
 description R3 port 0/1
 no switchport
 ip address 172.16.31.20 255.255.255.0
 ip access-group 100 out
!
access-list 100 permit ip 172.16.90.1 0.0.0.6 any
```

Note also that the ACL used does not permit any traffic other than that which is sourced from the test network because SW2 has no other Layer 3 segments connected but the test network. Any traffic leaving toward R3 is either from the test network or is sourced by SW2 itself. Traffic sourced by a router is not a subject for outbound inspection with security ACLs, so no additional permits need to be added into ACL 100.

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 Cisco IOS Software commands as well a collection of proprietary commands such as **show all**.

11. Multicast

Issue: Announce the shared root without use of any dense groups or static configurations.

Solution:

Because a shared root is involved, you know this configuration requirement is a sparse mode configuration. When you configure sparse mode, a rendezvous point is involved. The challenge with sparse mode is how to advertise to all sparse mode routers the location of the rendezvous

point. Three methods of advertising the rendezvous point are: static configuration, auto-rp using pim sparse-dense mode, and the multicast bootstrap routing protocol. The configuration requirement prohibits the use of "static configurations or dense groups." The static configuration restriction eliminates the static method of PIM sparse mode rendezvous point advertisement. The "dense group" restriction eliminates the auto-rp method because it requires PIM sparse/dense mode. Therefore, only one rendezvous point advertisement method remains: the multicast bootstrap routing protocol.

Multicast Diagram

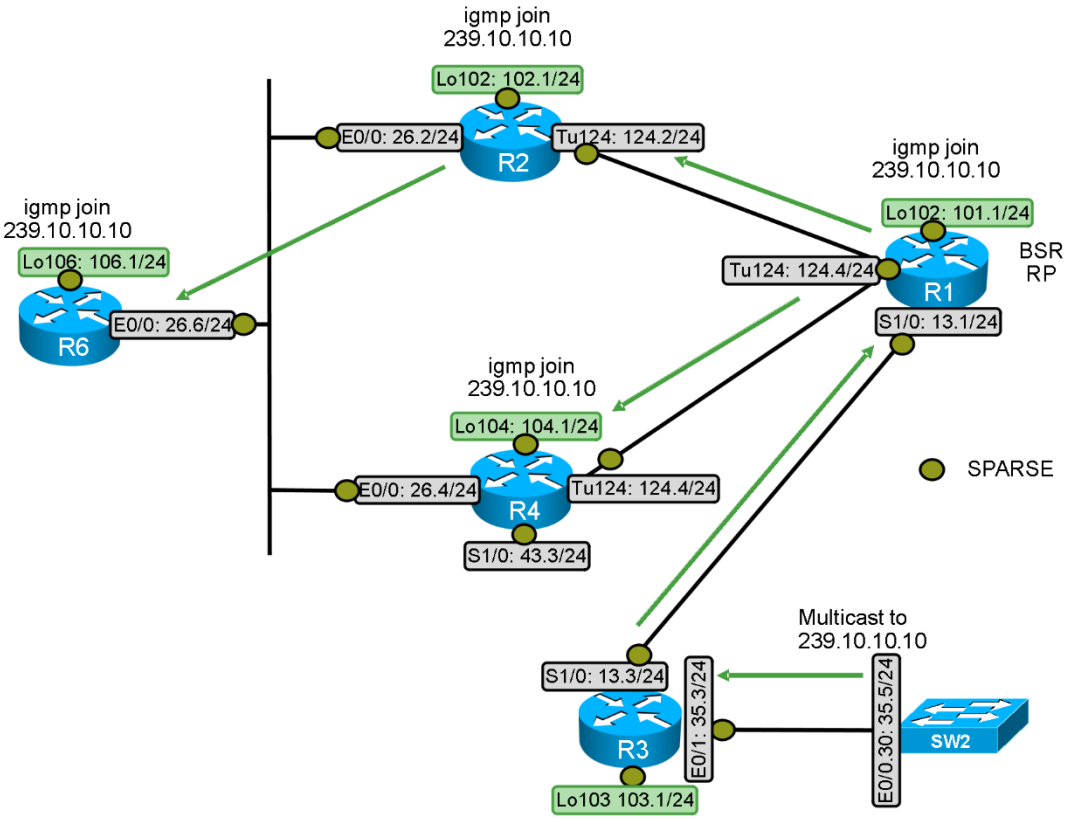

To fulfill this configuration requirement, configure PIM sparse mode with the Bootstrap Router (BSR) feature. Because R1 is specified to be configured as the shared root, configure R1 as the candidate rendezvous point for the 239.10.10.10 multicast group and for the bootstrap router.

Configure multicast on R1, R2, R3, R4, and R6.

R1:

```
interface Loopback101
 ip address 172.16.101.1 255.255.255.0
 ip pim sparse-mode
 ip igmp join-group 239.10.10.10
!
interface Tunnel124
 ip address 172.16.124.1 255.255.255.0
 no ip redirects
 ip pim sparse-mode
 ip nhrp map multicast 1.1.1.2
 ip nhrp map multicast 1.1.1.4
 ip nhrp network-id 10
```



```
 ip ospf network non-broadcast
 tunnel source Loopback1
 tunnel mode gre multipoint
 tunnel key 10
!
!
interface Serial1/0
 ip address 172.16.13.1 255.255.255.0
ip pim sparse-mode
!
ip pim bsr-candidate Loopback101 0
ip pim rp-candidate Loopback101 group-list 80
!
access-list 80 permit 239.10.10.10
!
```

R2:

```
interface Loopback102
 ip address 172.16.102.1 255.255.255.0
 ip pim sparse-mode
 ip igmp join-group 239.10.10.10
!
interface Tunnel124
 ip address 172.16.124.2 255.255.255.0
 no ip redirects
 ip pim sparse-mode
 ip nhrp map 172.16.124.1 1.1.1.1
 ip nhrp map multicast 1.1.1.1
 ip nhrp network-id 10
 ip nhrp nhs 172.16.124.1
 ip ospf network non-broadcast
 ip ospf priority 0
 tunnel source Loopback1
 tunnel mode gre multipoint
 tunnel key 10
!
interface Ethernet0/0
 ip address 172.16.26.2 255.255.255.0
 ip pim sparse-mode
!
```

R3:

```
interface Loopback103
 ip address 172.16.103.1 255.255.255.0
 ip pim sparse-mode
 ip igmp join-group 239.10.10.10
!
interface Ethernet0/1
 ip address 7.7.7.3 255.255.255.0 secondary
 ip address 172.16.31.3 255.255.255.0
 ip pim sparse-mode
!
!
interface Serial1/0
 ip address 172.16.13.3 255.255.255.0
 ip pim sparse-mode
!
```

R4:

```
interface Loopback104
 ip address 172.16.104.1 255.255.255.0
 ip pim sparse-mode
```

```
 ip igmp join-group 239.10.10.10
!
interface Tunnel124
 ip address 172.16.124.4 255.255.255.0
 no ip redirects
 ip pim sparse-mode
 ip nhrp map 172.16.124.1 1.1.1.1
 ip nhrp map multicast 1.1.1.1
 ip nhrp network-id 10
 ip nhrp nhs 172.16.124.1
 ip ospf network non-broadcast
 ip ospf priority 0
 tunnel source Loopback1
 tunnel mode gre multipoint
 tunnel key 10
!
!
interface Ethernet0/1
 ip address 172.16.26.4 255.255.255.0
 ip pim sparse-mode
!
```

R6:

```
interface Loopback106
 ip address 172.16.106.1 255.255.255.0
 ip pim sparse-mode
 ip igmp join-group 239.10.10.10
!
interface Ethernet0/0
 ip address 172.16.26.6 255.255.255.0
 ip pim sparse-mode
!
```

Use the **show ip pim** command to verify the BSR router on R1:

```
R1#show ip pim bsr-router
PIMv2 Bootstrap information
This system is the Bootstrap Router (BSR)
  BSR address: 172.16.101.1 (?)
  Uptime:      09:51:39, BSR Priority: 0, Hash mask length: 0
  Next bootstrap message in 00:00:25
  Candidate RP: 172.16.101.1(Loopback101)
    Holdtime 150 seconds
    Advertisement interval 60 seconds
    Next advertisement in 00:00:20
    Group acl: 80
R1#
```

Use the **show ip pim** command to verify the BSR router on R2, R3, R4, and R6:

```
R2#show ip pim bsr-router
PIMv2 Bootstrap information
  BSR address: 172.16.101.1 (?)
  Uptime:      09:53:24, BSR Priority: 0, Hash mask length: 0
  Expires:     00:01:53
R2#

R3#show ip pim bsr-router
PIMv2 Bootstrap information
  BSR address: 172.16.101.1 (?)
```



```
  Uptime:      09:52:29, BSR Priority: 0, Hash mask length: 0
  Expires:     00:01:46
R3#

R4#show ip pim bsr-router
PIMv2 Bootstrap information
  BSR address: 172.16.101.1 (?)
  Uptime:      09:51:47, BSR Priority: 0, Hash mask length: 0
  Expires:     00:01:27
R4#
R6#show ip pim bsr-router
PIMv2 Bootstrap information
  BSR address: 172.16.101.1 (?)
  Uptime:      09:50:03, BSR Priority: 0, Hash mask length: 0
  Expires:     00:01:11
R6#
```

Use the **ping** command to verify the ping from SW2:

```
SW2#ping 239.10.10.10 repeat 5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 239.10.10.10, timeout is 2 seconds:

Reply to request 0 from 172.16.103.1, 24 ms
Reply to request 0 from 172.16.106.1, 92 ms
Reply to request 0 from 172.16.104.1, 64 ms
Reply to request 0 from 172.16.102.1, 64 ms
Reply to request 0 from 172.16.101.1, 36 ms
Reply to request 1 from 172.16.103.1, 1 ms
Reply to request 1 from 172.16.106.1, 32 ms
Reply to request 1 from 172.16.102.1, 32 ms
Reply to request 1 from 172.16.104.1, 32 ms
Reply to request 1 from 172.16.101.1, 20 ms
Reply to request 1 from 172.16.106.1, 16 ms
Reply to request 1 from 172.16.102.1, 16 ms
Reply to request 1 from 172.16.104.1, 16 ms
Reply to request 1 from 172.16.101.1, 8 ms
Reply to request 2 from 172.16.103.1, 1 ms
Reply to request 2 from 172.16.106.1, 16 ms
Reply to request 2 from 172.16.102.1, 16 ms
Reply to request 2 from 172.16.104.1, 16 ms
Reply to request 2 from 172.16.101.1, 8 ms
Reply to request 3 from 172.16.103.1, 1 ms
Reply to request 3 from 172.16.106.1, 16 ms
Reply to request 3 from 172.16.102.1, 16 ms
Reply to request 3 from 172.16.104.1, 16 ms
Reply to request 3 from 172.16.101.1, 8 ms
Reply to request 4 from 172.16.103.1, 1 ms
Reply to request 4 from 172.16.106.1, 16 ms
Reply to request 4 from 172.16.102.1, 16 ms
Reply to request 4 from 172.16.104.1, 16 ms
Reply to request 4 from 172.16.101.1, 8 ms
SW2#
```

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 Cisco IOS Software commands as well a collection of proprietary commands such as **show all**.