360RSW05-WKB-CA08

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 8 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 8 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 as the requirements. You will *not* receive partial credit for partially completed subsections.
- IPv4 subnets that are displayed in the scenario diagram belong to network 160.20.0.0/16.
- IPv6 networks that are used in the scenario will use a FEC0::/9 network.
- *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 for IPv4 and IPv6 protocols.
- 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 involved in this scenario must be reachable, unless the instructions explicitly specify otherwise.

- Networks 192.50.*.* are excluded from the previous requirement to the extent explicitly described in this scenario.
- Unless the instructions explicitly specify otherwise, addresses and networks that advertised in the Border Gateway Protocol (BGP) section need to be reachable by all BGP routers but do not have to be reachable by routers that use only interior gateway protocol (IGP).
- 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.
- 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 160.20.0.0/16.
All IP addresses in this lab belong to the 160.20.0.0/16 address space with the exception of prefixes that are explicitly specified to be part of 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.

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 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," "Switch-to-Router Connections," and "Switch-to-Switch Connections" tables to analyze the VLAN propagation in this lab.

See the following diagram for the VLAN layout.

VLAN Propagation Diagram

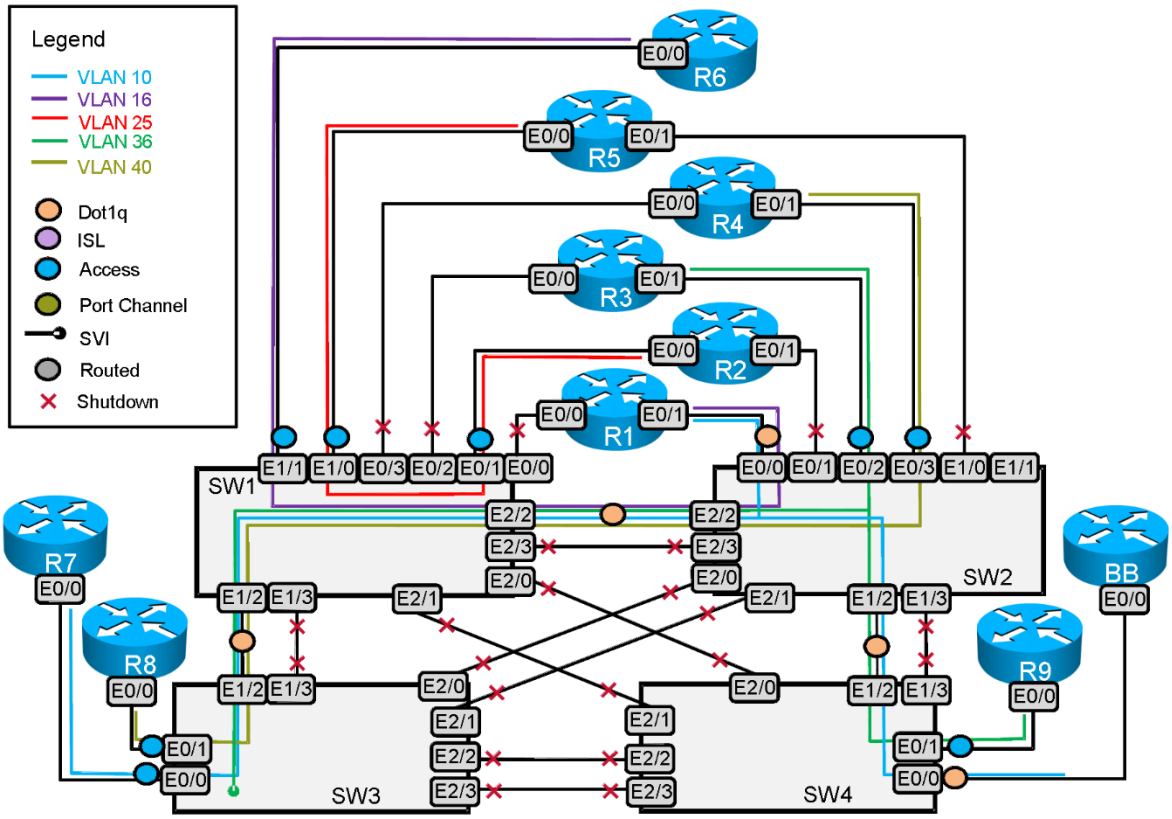

Carefully review the entire scenario. Closely examine the supplied diagram and any associated tables. Determine how you need to configure VTP, how to configure ports that are assigned as trunks, and how to configure ports that are assigned as static VLAN ports. Use the **switchport mode access** command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

Issue: On SW3 and SW4, create the VLANs as required by other sections of this scenario.

Solution:

Carefully examine which VLANs are used on SW3 and SW4, paying special attention to the "Quality of Service" section, which requires connectivity to an imaginary traffic generator on SW4 E0/0 VLAN 10. The VLANs that need to be configured on SW3 and SW4 are illustrated in the VLAN propagation diagram.

```
SW3#show vlan brie

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active    Et0/2, Et0/3, Et1/0, Et1/1
                                                Et1/3, Et2/0, Et2/1, Et2/2
                                                Et2/3
10   VLAN0010                         active    Et0/0
36   VLAN0036                         active
40   VLAN0040                         active    Et0/1
1002 fddi-default                     act/unsup
1003 token-ring-default               act/unsup
1004 fddinet-default                  act/unsup
1005 trnet-default                    act/unsup
SW3#
SW3#

SW4#show vlan brie

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active    Et0/0, Et0/2, Et0/3, Et1/0
                                                Et1/1, Et1/3, Et2/0, Et2/1
                                                Et2/2, Et2/3
10   VLAN0010                         active
36   VLAN0036                         active    Et0/1
1002 fddi-default                     act/unsup
1003 token-ring-default               act/unsup
1004 fddinet-default                  act/unsup
1005 trnet-default                    act/unsup
SW4#
```

Issue: Only VLANs that need to carry traffic between the switches should be allowed on the trunks.

Solution:

Examine the VLAN propagation diagram to establish which VLANs need to cross each of the interswitch trunks; note the differences in trunk encapsulation:

- E2/2 connecting SW1 and SW2 has dot1q encapsulation and permits VLANs 10, 16, 36, and 40.
- E1/2 connecting SW1 and SW3 has dot1q encapsulation and permits VLANs 10, 36, and 40.
- E1/2 connecting SW2 and SW4 has dot1q encapsulation and permits VLANs 10 and 36.

- SW4 E0/0 connects to the backbone with dot1q encapsulation and permits VLAN 10.

```
SW1#show int trunk

Port                Mode         Encapsulation  Status        Native vlan
Et1/2               on           802.1q         trunking      1
Et2/2               on           802.1q         trunking      1

Port                Vlans allowed on trunk
Et1/2               10,36,40
Et2/2               10,16,36,40

Port                Vlans allowed and active in management domain
Et1/2               10,36,40
Et2/2               10,16,36,40

Port                Vlans in spanning tree forwarding state and not pruned
Et1/2               10,36,40
Et2/2               10,16,36,40
SW1#

SW2#show int trunk

Port                Mode         Encapsulation  Status        Native vlan
Et0/0               on           802.1q         trunking      1
Et1/2               on           802.1q         trunking      1
Et2/2               on           802.1q         trunking      1

Port                Vlans allowed on trunk
Et0/0               10,16
Et1/2               10,36
Et2/2               10,16,36,40

Port                Vlans allowed and active in management domain
Et0/0               10,16
Et1/2               10,36
Et2/2               10,16,36,40

Port                Vlans in spanning tree forwarding state and not pruned
Et0/0               10,16
Et1/2               10,36
Et2/2               10,16,36,40
SW2#
SW3#show int trunk

Port                Mode         Encapsulation  Status        Native vlan
Et1/2               on           802.1q         trunking      1

Port                Vlans allowed on trunk
Et1/2               10,36,40

Port                Vlans allowed and active in management domain
Et1/2               10,36,40

Port                Vlans in spanning tree forwarding state and not pruned
Et1/2               10,36,40
SW3#

SW4#show int trunk

Port                Mode         Encapsulation  Status        Native vlan
Et0/0               on           802.1q         trunking      1
Et1/2               on           802.1q         trunking      1

Port                Vlans allowed on trunk
Et0/0               10
```

```
Et1/2               10,36

Port                Vlans allowed and active in management domain
Et0/0               10
Et1/2               10,36

Port                Vlans in spanning tree forwarding state and not pruned
Et0/0               none
Et1/2               10,36
SW4#
```

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 a collection of proprietary commands such as **show all**.

2. IPv4 OSPF

Issue: Configure OSPF on the 160.20.123.0/24 network between routers R1, R2, and R3. Use the OSPF nonbroadcast network type. Make the 160.20.123.0/24 network the OSPF backbone area. R1 should be the DR. R2 and R3 should be DROTHERs.

Solution:

A DR is elected to source the network link-state advertisement (LSA) and to reduce flooding of LSAs. Each router on the link must form an adjacency with the DR. To assure that R1 is elected as the DR, and R2 and R3 are DROTHERs, set the OSPF priority to 0 at the interface level on R2 and R3. The OSPF nonbroadcast network type requires neighbor statements so that OSPF packets can be unicast rather than multicast. At R1, enter two neighbor statements, one for R2 and the second for R3.

R1

```
interface Ethernet0/2
 ip address 160.20.123.1 255.255.255.0
 ip pim sparse-mode
 ip ospf network non-broadcast
!
router ospf 1
 network 160.20.123.0 0.0.0.255 area 0
 neighbor 160.20.123.2
 neighbor 160.20.123.3
!
```

R2

```
interface Ethernet0/2
 ip address 160.20.123.2 255.255.255.0
 ip pim sparse-mode
 ip ospf network non-broadcast
 ip ospf priority 0
!
router ospf 1
 network 160.20.123.0 0.0.0.255 area 0
!
```


R3

```
interface Ethernet0/2
 ip address 160.20.123.3 255.255.255.0
 ip ospf network non-broadcast
 ip ospf priority 0

!
router ospf 1
 network 160.20.123.0 0.0.0.255 area 0
!
```

Issue: On R3, place three loopback networks into OSPF Area 4. Summarize the three entries with the most efficient mask.

Solution:

With OSPF, you possess two summarization tools: the **area range** command and the **summary address** command. The **area range** command is used to summarize between OSPF areas and makes that router an Area Border Router (ABR). The **summary address** command is used to summarize routes from outside of OSPF, and configuring it makes that router an Autonomous System Boundary Router (ASBR). Since the prefixes to be summarized originate from an OSPF area, the **area range** command is the appropriate summarization tool. The most efficient mask is the one that summarizes the required addresses and as few others as possible. It is also the longest possible mask that will include the required networks. You can determine the most efficient mask by counting the number of bits, from left to right, that are identical in the specified network addresses. As you see below, the first 20 bits of each address are identical:

Dotted Decimal Network	Binary Network
160.20.163.0	10100000.00010100.10100011.00000000
160.20.169.0	10100000.00010100.10101001.00000000
160.20.174.0	10100000.00010100.10101110.00000000

R3

```
router ospf 1
 area 4 range 160.20.160.0 255.255.240.0

 network 160.20.123.0 0.0.0.255 area 0
 network 160.20.160.0 0.0.15.255 area 4
!
```

Issue: Configure OSPF Area 10 on the link between R1 and R6. Make sure that the routers in this OSPF area possess the minimum amount of routing information to reach all destinations within your pod.

Solution:

To minimize the amount of routing information in an OSPF area other than Area 0, configure the OSPF stub area feature. An OSPF stub area restricts external routes from entering the designated area. An OSPF totally stubby area restricts both external and interarea OSPF routes from entering

an area. When you configure an OSPF area other than Area 0 as a stub area, you need to configure all routers in the stub area with the **area stub** keywords, as well. If you want to configure an area as a totally stubby area, you must still configure all routers within the stub area with the **area X stub** command. On the router that is the ABR for the totally stubby area, enter the command **area X stub no-summary** under the OSPF routing process. Note that the ABR, R1 in this case, will advertise a default route into the stub area to provide connectivity outside the area.

R1:

```
router ospf 1
 area 10 stub no-summary

 network 160.20.16.0 0.0.0.255 area 10
 network 160.20.123.0 0.0.0.255 area 0
 neighbor 160.20.123.3
 neighbor 160.20.123.2
!
```

R6:

```
router ospf 1
 area 10 stub
 network 160.20.0.0 0.0.255.255 area 10
!
```

Issue: Configure OSPF Area 25 on the link between R2 and R5. Advertise the loopback 160.20.105.1/24 in OSPF Area 105.

Solution:

Since router R5 has no direct connection to Area 0, you will need to configure a virtual link through Area 25 to allow Area 105 to be accessible to the rest of the OSPF network. Be mindful of changes in the OSPF router ID in case new loopback interfaces are added or removed. It is recommended that you either finish creating loopbacks and assigning IP addresses before the virtual link is configured (or rebooting if the loopback addresses were changed) or manually enter an OSPF RID into the router configurations.

To verify that your virtual link is truly operational, look for a neighbor over the virtual link, and look for the string “adjacency state full” in the output of the **show ip ospf virtual-link** command:

R5:

```
router ospf 1
 router-id 160.20.105.1
 area 25 virtual-link 160.20.102.1
 network 160.20.25.0 0.0.0.255 area 25
 network 160.20.105.1 0.0.0.0 area 105


R5#show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
160.20.102.1      0   FULL/  -            -        160.20.25.2     OSPF_VL0
160.20.102.1      0   FULL/  -        00:00:32     160.20.25.2     Ethernet0/0
R5#
R5#show ip ospf virtual-links
Virtual Link OSPF_VL0 to router 160.20.102.1 is up
  Run as demand circuit
  DoNotAge LSA allowed.
  Transit area 25, via interface Ethernet0/0
 Topology-MTID    Cost    Disabled     Shutdown      Topology Name
        0          10        no          no            Base
  Transmit Delay is 1 sec, State POINT_TO_POINT,
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
```



```
    Hello due in 00:00:06
    Adjacency State FULL (Hello suppressed)
    Index 1/2, retransmission queue length 0, number of retransmission 0
    First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)
    Last retransmission scan length is 0, maximum is 0
    Last retransmission scan time is 0 msec, maximum is 0 msec
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 1000 Cisco IOS Software commands, as well a collection of proprietary commands such as **show all**.

3. IPv4 RIP

Issue: Configure a RIPv2 update exchange over the link between R1 and R4 and between R4 and R8. Make sure that RIP advertises only over the specified links.

Solution:

RIPv2 is enabled under the RIP routing process. You can control which interfaces RIP updates and gets advertised out of by using the **passive-interface** command. A recommended general practice to apply to a standard RIP configuration is to enter the **passive-interface default** command under the **router rip** configuration command and then enter **no passive-interface** *Xy*, where *Xy* is the specific interface you want RIP updates to be advertised on. You can verify both the RIP version and passive status with the **show ip protocols** command.

Example on R1:

```
router rip
 version 2
 passive-interface default
 no passive-interface Ethernet0/3
 network 160.20.0.0
 default-metric 1
 no auto-summary
!
```

Issue: On R8, advertise the Loopback192 interface networks into RIP without using a network statement.

Solution:

This task calls for redistribution of connected routes. You could use a route map referencing Loopback192 to include only the specified connected routes.

Issue: Make sure that R8 advertises only the following networks to R4; use the minimum number of ACL statements to match these networks.

- 192.50.153.0
- 192.50.155.0
- 192.50.157.0
- 192.50.159.0

Solution:

The solution uses a distribute list under the RIP process applied outbound from interface VLAN 40. As shown, these network addresses can be matched using a single-line ACL:

Dotted Decimal Network	Binary Network
192.50.153.0	11000000.00110010.10011**00**1.0000
192.50.155.0	11000000.00110010.10011**01**1.0000
192.50.157.0	11000000.00110010.10011**10**1.0000
192.50.159.0	11000000.00110010.10011**11**1.0000

Notice that these networks differ only in bits 6 and 7 of the third octet. You can therefore match these four networks by allowing just these two bits to vary in the inverse mask: 0.0.6.0.

Here is an example on R8:

```
router rip
 version 2
 redistribute connected
 passive-interface default
 no passive-interface Ethernet0/0
 network 160.20.0.0
 default-metric 1
 distribute-list FILTER192 out Ethernet0/0
 no auto-summary
!
!
ip access-list standard FILTER192
 permit 192.50.153.0 0.0.6.0
!
```

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 a collection of proprietary commands such as **show all**.

4. IPv4 EIGRP

Issue: Make sure that R8 advertises only the following networks to R7. Use the minimum number of ACL statements to match these networks.

- 192.50.153.0/24
- 192.50.155.0/24
- 192.50.157.0/24
- 192.50.159.0/24

Solution:

This is similar to the filtering requirement in the RIP section, and the same ACL can be used. Apply it using a distribute list under the EIGRP 10 process on R9, outbound from Interface E0/1.

Here is an example on R8:

```
router eigrp 10
```



```
 distribute-list FILTER192 out Ethernet0/1
 default-metric 1544 2000 255 1 1500
 network 160.20.17.0 0.0.0.255
 redistribute connected
!
!
ip access-list standard FILTER192
 permit 192.50.153.0 0.0.6.0
!
```

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 a collection of proprietary commands such as **show all**.

5. 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, 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 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

Redistribution Point	Into RIP		Into OSPF		Into EIGRP 10	
	Permit	**Deny**	**Permit**	**Deny**	**Permit**	**Deny**
R1			All EIGRP routes All RIP routes		All OSPF routes All RIP routes	
R3	All OSPF routes		All RIP routes			
R7	Selected connected	Selected connected			Selected connected	Selected connected

Note that neither OSPF nor EIGRP is redistributed into RIP on R1. How will R4 obtain a route to prefixes in these domains? Both policy routing and default routes are forbidden. The solution that is used in the Mentor Guide takes advantage of the fact that all of the missing routes are subnets of 160.20.0.0/16. R1 is advertising some prefixes from this range to R4 using RIP. The following command implemented on R1 E0/3 will cause the RIP process on R1 to advertise 160.20.0.0/16 to R4:

```
ip summary-address rip 160.20.0.0 255.255.0.0
```

On Cisco routers, you can use a Tool Command Language (Tcl) script like the one below to automate reachability testing.

```
tclsh
foreach address {

160.20.16.1
160.20.14.1
160.20.10.1
160.20.123.1
160.20.101.1

160.20.25.2
160.20.123.2
160.20.102.1

160.20.36.130
160.20.163.1
160.20.174.1
160.20.169.1
160.20.36.3
160.20.123.3
160.20.103.1

160.20.40.4
160.20.14.4
160.20.104.1

160.20.25.5
160.20.105.1

5.5.5.1
4.4.4.1
160.20.16.6
160.20.106.1

160.20.17.10
160.20.10.10

192.50.153.1
192.50.155.1
192.50.157.1
192.50.159.1

160.20.40.20
160.20.17.20

7.5.5.1
7.4.4.1

160.20.36.6
160.20.107.1

} {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 a collection of proprietary commands such as **show all**.

6. Border Gateway Protocol

BGP Diagram

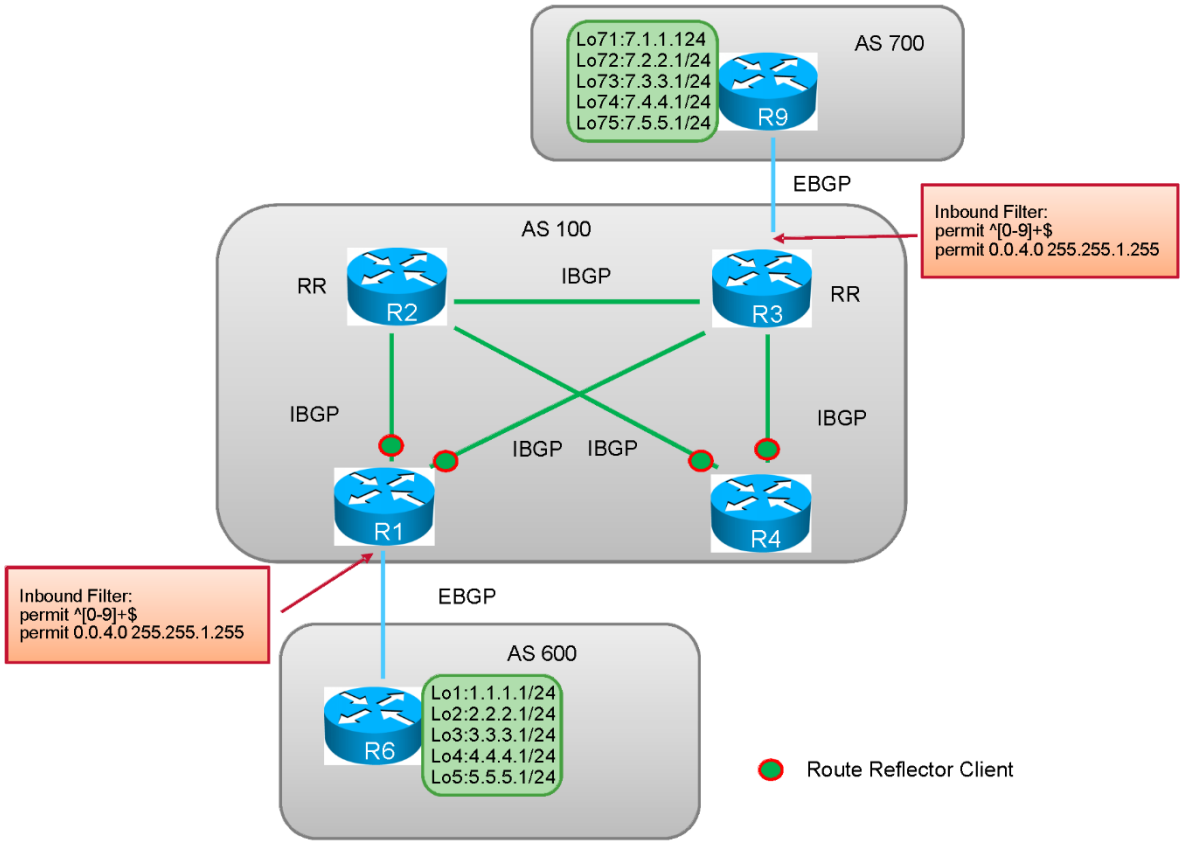

Issue: Make R1, R2, R3, and R4 BGP speakers within AS 100. Do not allow a full mesh of Interior Border Gateway Protocol (IBGP) speakers within AS 100. Provide redundant Network Layer Reachability Information (NLRI) exchange using routers R2 and R3.

Solution:

You have three options when configuring IBGP neighbor relationships:

1. Configure a full mesh of IBGP speakers.
2. Configure a route reflector or a collection of route reflectors within an autonomous system (AS).
3. Configure a confederation of private autonomous systems.

Route reflectors can be combined with confederations. In this configuration, you are instructed not to allow a full mesh of IBGP speakers within AS 10. This requirement can be fulfilled with either a route reflector or confederation. However, the requirement to provide redundant NLRI

exchange suggests route reflectors as the solution. A redundant route reflector can provide the redundant NLRI exchange.

Both R2 and R3 will be configured as route reflectors with the remaining routers in AS 10 being route reflector clients to both R2 and R3. Both routers R2 and R3 will be configured with an identical BGP cluster ID to avoid loops. Remember that you must enter the BGP cluster ID before you enter the **route-reflector-client** commands.

On R2:

```
R2#show run | begin router bgp
router bgp 100
 bgp cluster-id 1
 bgp log-neighbor-changes
 neighbor 160.20.101.1 remote-as 100
 neighbor 160.20.101.1 update-source Loopback102
 neighbor 160.20.101.1 route-reflector-client
 neighbor 160.20.103.1 remote-as 100
 neighbor 160.20.103.1 update-source Loopback102
 neighbor 160.20.104.1 remote-as 100
 neighbor 160.20.104.1 update-source Loopback102
 neighbor 160.20.104.1 route-reflector-client
 no auto-summary
!
```

On R3:

```
R3#show run | begin router bgp
router bgp 100
 bgp cluster-id 1
 bgp log-neighbor-changes
 neighbor 160.20.36.6 remote-as 700
 neighbor 160.20.36.6 route-map eBGP-in in
 neighbor 160.20.101.1 remote-as 100
 neighbor 160.20.101.1 update-source Loopback103
 neighbor 160.20.101.1 route-reflector-client
 neighbor 160.20.102.1 remote-as 100
 neighbor 160.20.102.1 update-source Loopback103
 neighbor 160.20.104.1 remote-as 100
 neighbor 160.20.104.1 update-source Loopback103
 neighbor 160.20.104.1 route-reflector-client
 no auto-summary
!
```

Issue: Configure R1 and R3 so that AS 100 only accepts from its External Border Gateway Protocol (EBGP) peers those prefixes with a third octet of 4 or 5 that originated from the connected AS.

Solution:

This filtering requirement possesses two components, one based upon an IP prefix and one based upon AS-path characteristics. However, in order for a BGP update to be accepted, both criteria must match, resulting in a logical AND match requirement. A logical AND match requirement is best fulfilled by configuring a route map. Configure a route map with the two match statements under a single route map stanza: one match criterion for a prefix match and a second match criterion based upon an AS-path match. The route map is applied to the inbound EBGP neighbor statement on R1 and R3. The supporting access lists for the route map could be written as follows.

For the prefix match: **access-list 1 permit 0.0.4.0 255.255.1.255**

This combination of base and mask says that the first, second, and fourth octets can be anything, but the third octet must be 0000 0100 or 0000 0101 (4 or 5). The regular expression below says

that the path must contain only one or more (+) numerals. This expression does not permit a blank path or a path with any spaces. Since there must be at least one numeral and there can be no spaces, the path must consist of a single AS number, which would be true only of prefixes that originated from neighboring autonomous systems. If these filters are working as designed, then you should see only the prefixes 4.4.4.0, 5.5.5.0, 7.4.4.0, and 7.5.5.0 in AS 100.

For the AS-path match: **ip as-path access-list 1 permit ^[0-9]+$**

On R1 and R3:

Router BGP configuration

```
neighbor 160.20.36.6 route-map eBGP-in in
```

Router global configuration

```
ip access-list standard eBGP-in
 permit 0.0.4.0 255.255.1.255
!
ip as-path access-list 1 permit ^[0-9]+$
!
route-map eBGP-in permit 10
 match ip address eBGP-in
 match as-path 1
```

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 a collection of proprietary commands such as **show all**.

7. Router Maintenance

Issue: Ensure that workstations on the 11.1.1.0/24 private address space connected to VLAN 36 can reach the rest of the network using a portion of the address space of the R3 160.20.36.0/22 subnet. SW3 should be reachable from the rest of the network using the 160.20.36.130 address.

Solution:

This Network Address Translation (NAT) problem does not explicitly mention NAT. The router performing NAT—in this scenario it is router R3—will rewrite the source IP address of packets that originate from the 11.1.1.0/24 subnet to a source IP address that is known to router R1. Router R1 will never know that the packets originated from the 11.1.1.0/24 subnet.

The NAT inside interface is the Ethernet interface attached to VLAN 36 on R3, and the NAT outside interface is the R3 Ethernet0/2 interface. Line 1 below defines the inside source addresses that can be translated. Line 2 creates a pool of global addresses called **mypool**. This pool consists of a portion of the connected 160.20.36.0/22 network, which is advertised via OSPF to the rest of the pod. Line 3 creates the translation from permitted inside addresses to the global pool. Though not strictly necessary, given this topology, you could use the **overload** keyword because you are permitting up to 254 inside addresses to be translated to just half that number of global addresses. Finally, the **ip nat inside source static 11.1.1.7 160.20.36.130** command makes the SW3 private address reachable through the global address 160.20.36.130.

R3

```
interface Ethernet0/2
 ip address 160.20.123.3 255.255.255.0
```

```
 ip nat outside
 ip virtual-reassembly
!
interface Ethernet0/1
 ip address 11.1.1.3 255.255.255.0 secondary
 ip address 160.20.36.3 255.255.252.0
 ip nat inside
!
ip nat pool mypool 160.20.36.131 160.20.36.254 netmask 255.255.252.0
ip nat inside source list 11 pool mypool overload
ip nat inside source static 11.1.1.7 160.20.36.130
!
!
access-list 11 permit 11.1.1.0 0.0.0.255
```

Verification:

Ping from subnet 11.1.1.0/24 of R3:

```
R3#ping 160.20.16.6 source 11.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 160.20.16.6, timeout is 2 seconds:
Packet sent with a source address of 11.1.1.3
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 48/50/52 ms
R3#sh ip nat translations
Pro Inside global      Inside local       Outside local      Outside global
icmp 160.20.36.131:40  11.1.1.3:40        160.20.16.6:40     160.20.16.6:40
--- 160.20.36.130      11.1.1.7           ---                ---
R3#
```

Ping SW3 using 160.20.36.130:

```
R2#ping 160.20.36.130

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 160.20.36.130, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 100/101/108 ms
R2#

R3#show ip nat translations
Pro Inside global      Inside local       Outside local      Outside global
icmp 160.20.36.130:39  11.1.1.7:39        160.20.123.2:39    160.20.123.2:39
--- 160.20.36.130      11.1.1.7           ---                ---
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 a collection of proprietary commands such as **show all**.

8. IPv6 Routing

Issue: Configure IPv6 addresses.

Solution:

On each router, enter the command **ipv6 unicast-routing** in global configuration mode. Then assign the addresses as directed. See the Mentor Guide engine for specific configuration.

IPv6 configuration on an interface includes setting the IPv6 address using **ipv6 address** command. Additional configuration must be done on NBMA networks, such as assigning map statements and possibly link-local addresses for easy administration.

R2 (example):

```
interface Ethernet0/0
 ip address 160.20.25.2 255.255.255.0
 ipv6 address FE80::252 link-local
 ipv6 address FEC0:25::2/64
!
```

R5 (example):

```
interface Loopback105
 ip address 160.20.105.1 255.255.255.0
 ipv6 address FEC0:105::1/64
!
interface Ethernet0/0
 ip address 160.20.25.5 255.255.255.0
 ipv6 address FE80::255 link-local
 ipv6 address FEC0:25::5/64
!
```

Issue: Configure a logical point-to-point link between R2 and R6 to carry IPv6. Do not enable IPv6 on R1.

Solution:

This task requires configuration of a manual tunnel between R2 and R6; either IPv6 IP or GRE can be used as the encapsulation protocol. The IPv6 IP manual tunnel can carry only IPv6 packets, whereas the GRE encapsulated tunnel can carry many types of traffic.

R2

```
interface Tunnel10
 no ip address
 ipv6 address FEC0:26::2/64
 tunnel source Ethernet0/2
 tunnel destination 160.20.16.6
 tunnel mode ipv6ip
!
```

R6

```
interface Tunnel10
 no ip address
 ipv6 address FEC0:26::6/64
 tunnel source Ethernet0/0
 tunnel destination 160.20.123.2
 tunnel mode ipv6ip
!
```

Issue: Configure OSPFv3 Area 25 on the FEC0:25::/64 subnet. Use the OSPFv3 NBMA network.

Solution:

Enable OSPFv3 on each specified interface with the command **ipv6 ospf 1 area X**. The OSPF network type on the interfaces connecting routers R2 and R5 is NBMA, so you will need to configure neighbor statements on these interfaces:

R2

```
    interface Ethernet0/0
 ip address 160.20.25.2 255.255.255.0
 ipv6 address FE80::252 link-local
 ipv6 address FEC0:25::2/64
 ipv6 ospf network non-broadcast
 ipv6 ospf neighbor FE80::255
 ipv6 ospf 1 area 25
!
```

R5

```
interface Ethernet0/0
 ip address 160.20.25.5 255.255.255.0
 ipv6 address FE80::255 link-local
 ipv6 address FEC0:25::5/64
 ipv6 ospf network non-broadcast
 ipv6 ospf 1 area 25
!
```

Issue: Make sure that routers R2 and R5 see only a summary route representing the IPv6 loopback interfaces on R6

Solution:

This task can be accomplished by configuring an OSPFv3 interarea summary on R6. The challenges lie in recognizing that the addresses are given in hexadecimal and finding the appropriate mask length. The original addresses were configured with /124 masks. The last 32 bits of each address can be analyzed as follows:

```
::250:1       ::0000 0010 0101 0000:0000:0000:0000:0001
::251:1       ::0000 0010 0101 0001:0000:0000:0000:0001
::252:1       ::0000 0010 0110 0010:0000:0000:0000:0001
```

With the /124 mask, the host bits are those marked in green. The most efficient summary would be a /110, which is the number of bits in common. The command **area 6 range FEC0::250:0/110** under the **ipv6 router ospf** process on R6 was issued. Here is the resulting routing table on R2:

```
R2#show ipv6 route ospf
OI  FEC0::250:0/110 [110/11111]
     via FE80::A014:1006, Tunnel0
O   FEC0:105::1/128 [110/1]
     via FE80::255, Ethernet0/0
```

Issue: All traffic that leaves R6 destined to IPv6 addresses should be tagged with the DSCP value Expedited Forwarding (EF).

Solution:

The simplest solution is to take advantage of the 12.2(8)T Tunnel type of service (ToS) feature. Entering the command **tunnel tos 184** on R6 sets the ToS byte to binary 10111000 in the outer

(IPv4) header, regardless of the value of the ToS byte in the encapsulated IPv6 packet. Remember that DSCP is represented by the first six bits of the ToS byte, so 101110 is decimal 46 and Expedited Forwarding Per-Hop Behavior (EF PHB), but the command sets the full 8 bits and must be padded with two zeros on the right, resulting in decimal 184 (multiplication of the DSCP decimal value by 4 is an equivalent of padding with two zeros on the right and is the easiest way to calculate ToS value). This feature is documented at the following link.

http://www.cisco.com/en/US/docs/ios/12_0s/feature/guide/12s_tos.html

R6:

```
interface Tunnel10
 no ip address
 ipv6 address FEC0:26::6/64
 ipv6 ospf 1 area 0
 tunnel source Ethernet0/0
 tunnel destination 160.20.123.2
 tunnel mode ipv6ip
 tunnel tos 184
!
```

Issue: Configure R1 to monitor traffic (as described in the previous task) in transit from R6 to R2; count separately the number of R6-to-R2 IPv6 packets marked with the DSCP value EF and the number of R6-to-R2 IPv6 packets with all other DSCP values.
Configure R2 to monitor how many IPv6 packets have arrived from R6; count separately the number of packets marked with the DSCP value EF and the number of packets with all other DSCP values.

Solution:

To verify the operation of marking described in the previous task, access lists were created that matched on DSCP EF and were applied inbound on the Et0/1.16 subinterface of R1 and the Tunnel interface of R2. When you generated pings from R6 to FEC0:105::1, you get hits on R1 (IPv4), but not on R2 (IPv6).

R1

```
interface Ethernet0/1.16
 description VLAN16
 encapsulation dot1Q 16
 ip address 160.20.16.1 255.255.255.0
 ip access-group EF_counter in
 ip pim sparse-mode
 no snmp trap link-status
!
ip access-list extended EF_counter
 permit 41 host 160.20.16.6 host 160.20.123.2 dscp ef
 permit 41 host 160.20.16.6 host 160.20.123.2
 permit ip any any
!

R1#show access-lists EF_counter
Extended IP access list EF_counter
    10 permit 41 host 160.20.16.6 host 160.20.123.2 dscp ef (1141 matches)
    20 permit 41 host 160.20.16.6 host 160.20.123.2
    30 permit ip any any (3304 matches)
R1#
```

R2

```
interface Tunnel0
 no ip address
```

```
 ipv6 address FEC0:26::2/64
 ipv6 traffic-filter NOEF in
 ipv6 ospf 1 area 0
 tunnel source Ethernet0/2
 tunnel destination 160.20.16.6
 tunnel mode ipv6ip
!
ipv6 access-list NOEF
 permit ipv6 any any dscp ef
 permit ipv6 any any
!
R2#show access-lists NOEF
IPv6 access list NOEF
    permit ipv6 any any dscp ef sequence 10
    permit ipv6 any any (1152 matches) sequence 20
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 1000 Cisco IOS Software commands, as well a collection of proprietary commands such as **show all**.

9. Quality of Service

Issue: Limit incoming UDP traffic that is destined to port 5111 to a rate of 8000 b/s on the R1 interface on subnet 160.20.10.0/24. Configure the minimum values for burst size and extended burst size. The solution should continue to provide effective limiting if the packet size on the traffic generator is changed. Drop excessive traffic.

Solution:

Even though the term “traffic policing” is never explicitly stated, this task possesses all the characteristics of a traffic-policing requirement. An inbound rate is specified for a defined set of traffic. If that inbound rate is exceeded, the traffic is to be dropped. Traffic policing will never buffer traffic. A traffic-policing configuration will drop traffic, pass it through, or pass it through and mark the packets using tools such as IP precedence bits or DSCP bits. This configuration requirement can be fulfilled using the **rate-limit** command (committed access rate [CAR]) or the Modular QoS CLI (MQC).

When completing this task, pay attention to the specified packet size. Depending on the policer that is used, the burst byte count can be as small as 1000 bytes. However, specifying such a low burst size will make packets always exceed the burst because the traffic generator packet size is 1024 bytes. The additional requirement to continue to provide effective limiting if the packet size changes implies that packet size can vary up to 1500 bytes (Ethernet IP MTU). Consequently, you must set up the policer to use 1500 as a normal burst value. The excess burst size needs to be configured to the minimum value that is allowable by Cisco IOS Software syntax.

The following example uses CAR to configure the policing:

R1

```
interface Ethernet0/1.10
 encapsulation dot1Q 10
 ip address 160.20.10.1 255.255.255.0
 rate-limit input access-group 107 8000 1500 2000 conform-action transmit exceed-
action drop
!
access-list 107 remark -------- CAR rate limited --------
access-list 107 permit udp any host 160.20.10.1 eq 5111
access-list 107 remark ---------------------------------
!
```

You can verify your configuration with the commands **show interface rate-limit** (CAR configuration) or **show policy-map interface** (MQC configuration).

```
R1#show interface rate-limit
Ethernet0/1.10 VLAN10
  Input
    matches: access-group 107
      params:  8000 bps, 1500 limit, 2000 extended limit
      conformed 2 packets, 156 bytes; action: transmit
      exceeded 0 packets, 0 bytes; action: drop
      last packet: 332770936ms ago, current burst: 0 bytes
      last cleared 5d23h ago, conformed 0 bps, exceeded 0 bps
```

10. Network Security

Issue: Configure R2 as a Secure Copy Protocol (SCP) server for the locally configured user noc with the password **cisco**.

Solution:

SCP uses the Secure Shell (SSH) protocol to permit remote copying of local files in a secure manner. Documentation for this feature can be found here:

http://www.cisco.com/en/US/docs/ios/sec_user_services/configuration/guide/sec_secure_copy_ps6922_TSD_Products_Configuration_Guide_Chapter.html

Here are the required steps:

1. Configure a domain name on the router. This configuration is required to complete the next step.

```
R2#show run | inc testlab
ip domain name testlab.com
R2#
```

2. Enter the command **crypto key generate** to create a set of keys to use for encryption. RSA and a 1024-bit key are used in this example:

```
R2#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R2(config)#crypto key generate rsa
The name for the keys will be: R2.testlab.com
Choose the size of the key modulus in the range of 360 to 4096 for your
  General Purpose Keys. Choosing a key modulus greater than 512 may take
  a few minutes.

How many bits in the modulus [512]: 1024
% Generating 1024 bit RSA keys, keys will be non-exportable...
[OK] (elapsed time was 0 seconds)

R2(config)#
*Jun 18 15:29:55.507: %SSH-5-ENABLED: SSH 1.99 has been enabled
R2(config)#end
R2#
```

3. Since you do not have a TACACS+ or RADIUS server available, you do not need to enable AAA. Instead, create a local user database, making sure to specify privilege level 15: **username noc privilege 15 password 0 cisco.**
4. Enter the command **ip scp server enable**.
5. By default, vty lines permit only Telnet connections. On vty lines 0 to 4, enter the commands **transport input ssh** and **login local**.

```
R2#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R2(config)#line vty 0 4
R2(config-line)#transport input ssh
R2(config-line)#login local
R2(config-line)#end
R2#
```

6. To verify the configuration, issue the following command from R5:

```
R5#more scp://noc@160.20.25.2/running-config
Password:

!
! Last configuration change at 07:32:40 PST Tue Jun 18 2013
version 15.3
```



```
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
!
hostname R2
!
boot-start-marker
boot-end-marker
!
<skipped>
```

This configuration should permit you to read the running configuration on R2 from R5 in a secure manner. If desired, you could also copy any file on R2 using the **copy scp:** command. Note that if you fail to authorize the user at privilege level 15, you will see output similar to the following:

```
R5#
R5#more scp://noc@160.20.25.2/running-config
Password:

Password:

% Authentication failed.

%Error opening scp://*****@160.20.25.2/running-config (Permission denied)
R5#
```

Issue: Configure R2 vty lines 0 to 4 to permit only SSH connections from IP address 160.20.25.5. Deny all other connections to the vty lines.

Solution:

To restrict source IP addresses, apply an access class to the specified vty lines. To deny all other connections to the router, enter the command **transport input none** on the remaining vty lines, if any. Configuration details can be seen in the Mentor Guide.

```
ip access-list standard SSH
 permit 160.20.25.5
!
line vty 0 4
 location cierswbv5-ce-lab08-sc, SJ
 access-class SSH in
 exec-timeout 0 0
 privilege level 15
 login local
 transport input ssh
!
```

Though you can no longer connect via Telnet to R2, you could use SSH to remotely manage it from R5 with this command:

```
ssh -l noc 160.20.25.2
```

11. Switch Specialties

Issue: Set port 1/1 of SW2 to bypass the learning and listening states of spanning tree.

Solution:

To bypass the learning and listening states of spanning tree, configure the following port configuration command on the switch: **spanning-tree portfast**.

```
interface Ethernet1/1
 description Bypass STP listen/learn
 spanning-tree portfast
```

Issue: Configure SW2 so that learned MAC addresses in VLAN 40 are retained for a period that is 1.5 times as long as the default.

Solution:

The default aging time for a MAC address residing in a MAC address table on a switch is 300 seconds. To fulfill this requirement, enter the following command in global configuration mode on SW2:

```
mac address-table aging-time 450 vlan 40
```

12. Multicast

Issue: Enable a multicast routing protocol that will use any unicast routing protocol for source address determination and that is also based on a shared tree. Do not use any dynamic methods to discover or advertise the root of the shared tree.

Solution:

The multicast routing protocol that meets the requirements stated above is Protocol Independent Multicast sparse mode (PIM-SM). If you cannot use any dynamic methods to discover or advertise the root of the shared tree, you must configure each multicast router with a manual configuration that statically identifies the rendezvous point (RP). This can be accomplished with the global configuration command **ip pim rp-address** *X.X.X.X*, where *X.X.X.X* is the IP address of the RP.

Example on R3:

```
R3#show run | inc ip pim rp-address
ip pim rp-address 160.20.101.1
R3#
```

Issue: Configure R1, R2, R3, R5, and R6 to join the multicast group 227.7.7.7. Associate this group with a loopback interface on each router.

Here is an example on R3:

```
interface Loopback103
 ip address 160.20.103.1 255.255.255.0
 ip pim sparse-mode
 ip igmp join-group 227.7.7.7
```


Solution:

In this lab, you are using the loopback interfaces to simulate subnets with multicast clients attached. You can simulate a multicast client by configuring a multicast routing protocol and using an **ip igmp join-group** command on the interface.

You should get replies from your pings similar to the following:

```
R4#ping 227.7.7.7 source 160.20.14.4

Type escape sequence to abort.
Sending 1, 100-byte ICMP Echos to 227.7.7.7, timeout is 2 seconds:
Packet sent with a source address of 160.20.14.4

Reply to request 0 from 160.20.101.1, 18 ms
Reply to request 0 from 160.20.105.1, 53 ms
Reply to request 0 from 160.20.102.1, 40 ms
Reply to request 0 from 160.20.106.1, 38 ms
Reply to request 0 from 160.20.103.1, 26 ms
R4#
```

Multicast Diagram

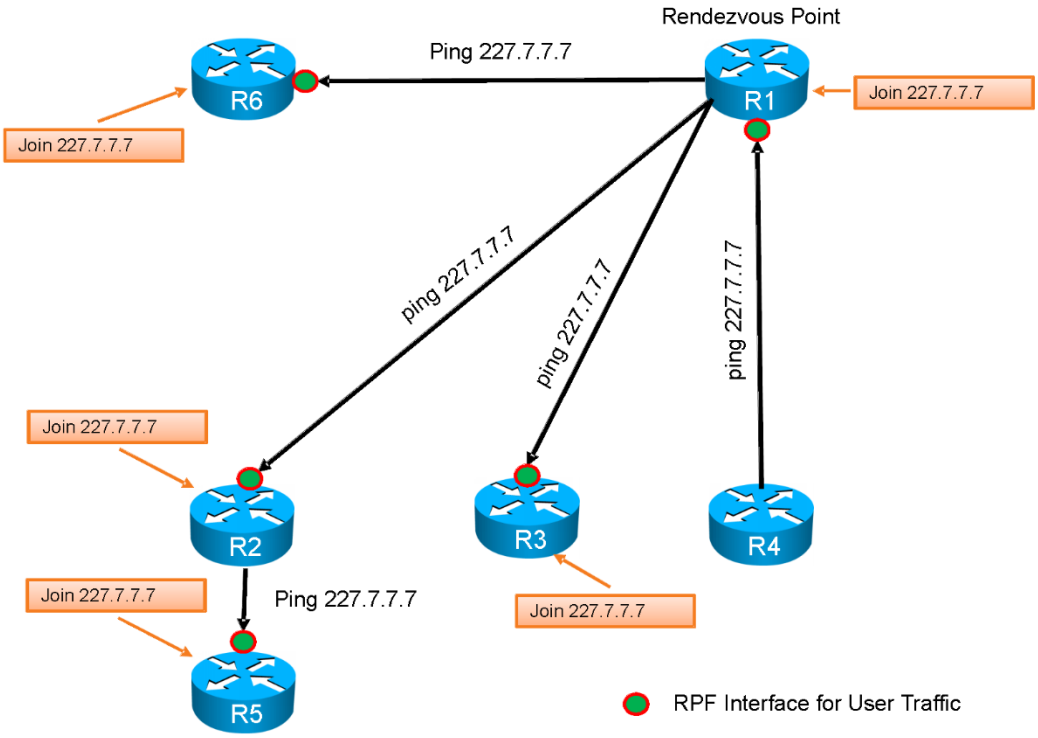

For a more detailed look at the multicast performance, examine the output of **show ip mroute**, as shown below for R2. Make sure that you have an active ping going from R4, because all these states will time out in a few minutes.

```
R2#show ip mroute
<skipped>
 (*, 227.7.7.7), 01:52:50/00:02:44, RP 160.20.101.1, flags: SJCL
  Incoming interface: Ethernet0/2, RPF nbr 160.20.123.1
  Outgoing interface list:
    Ethernet0/0, Forward/Sparse, 01:51:14/00:02:44
    Loopback102, Forward/Sparse, 01:52:50/00:02:24
 (160.20.14.4, 227.7.7.7), 00:01:03/00:03:28, flags: LT
  Incoming interface: Ethernet0/2, RPF nbr 160.20.123.1
  Outgoing interface list:
    Loopback102, Forward/Sparse, 00:01:04/00:02:23
    Ethernet0/0, Forward/Sparse, 00:01:04/00:03:24
<skipped>
```

The (*,G) entry above has an "S" flag, indicating that this group is being distributed in sparse mode. This entry describes the RPF interface and the outgoing interface list (OIL) for the shared tree, which is rooted at the RP. The (S,G) entry below the (*,G) entry shows the RPF interface and OIL for the source tree, rooted at the source of your ping.

Issue: Configure R3 so a workstation on VLAN 36 can only join multicast group 227.7.7.7 and cannot join any other groups.

Solution:

To restrict the groups that clients attached to a router interface may join, you can define permitted groups in an access list and apply it to the interface with the **ip igmp access-group** command on R3:

```
interface Ethernet0/1
 ip address 11.1.1.3 255.255.255.0 secondary
 ip address 160.20.36.3 255.255.252.0
 ip pim sparse-mode
 ip igmp access-group Multicast-Limit-IGMP-joins
!
ip access-list standard Multicast-Limit-IGMP-joins
 permit 227.7.7.7
!
```

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 a collection of proprietary commands such as **show all**.

