360RSW05-WKB-CA03

Cisco 360 CCIE R&S Exercise Workbook Introduction

The Cisco 360 CCIE® R&S Exercise Workbook contains 20 challenging scenarios at the Cisco 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 3 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 3 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 135.15.0.0/16.
- *Points will be deducted from multiple sections for failing to assign correct IPv4 addresses.*
- IPv6 subnets that are displayed in the scenario diagram belong to network FEC2::/16.
- *Points will be deducted from multiple sections for failing to assign correct IPv6 addresses.*
- Do not use any static routes unless explicitly permitted.
- Advertise loopback interfaces with their original masks.
- IPv4 network 0.0.0.0/0 should not appear in any routing table (**show ip route**).
- IPv6 network ::/0 should not appear in any routing table (**show ipv6 route**) except on R5.
- Do not introduce any new IP addresses.

- All IP addresses that are involved in this scenario must be reachable, unless explicitly specified otherwise.
- 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 have to be reachable by routers that use only IGP.
- Except in the Traffic Control subsection, use conventional routing algorithms only, unless specified otherwise.
- Do not create new interfaces to fulfill IGP requirements, and do not summarize unless you are explicitly asked to do so.
- 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 135.15.0.0/16.

All IP addresses in this lab belong to the 135.15.0.0/16 address space, except for prefixes that are used in the BGP 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 routers that use only IGP.

This statement relaxes the requirement that all IP addresses must be reachable. The BGP prefixes only need to 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 must be able to forward traffic to the addresses that are 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, the following 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. 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 seven 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. VLAN names are case-sensitive.

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

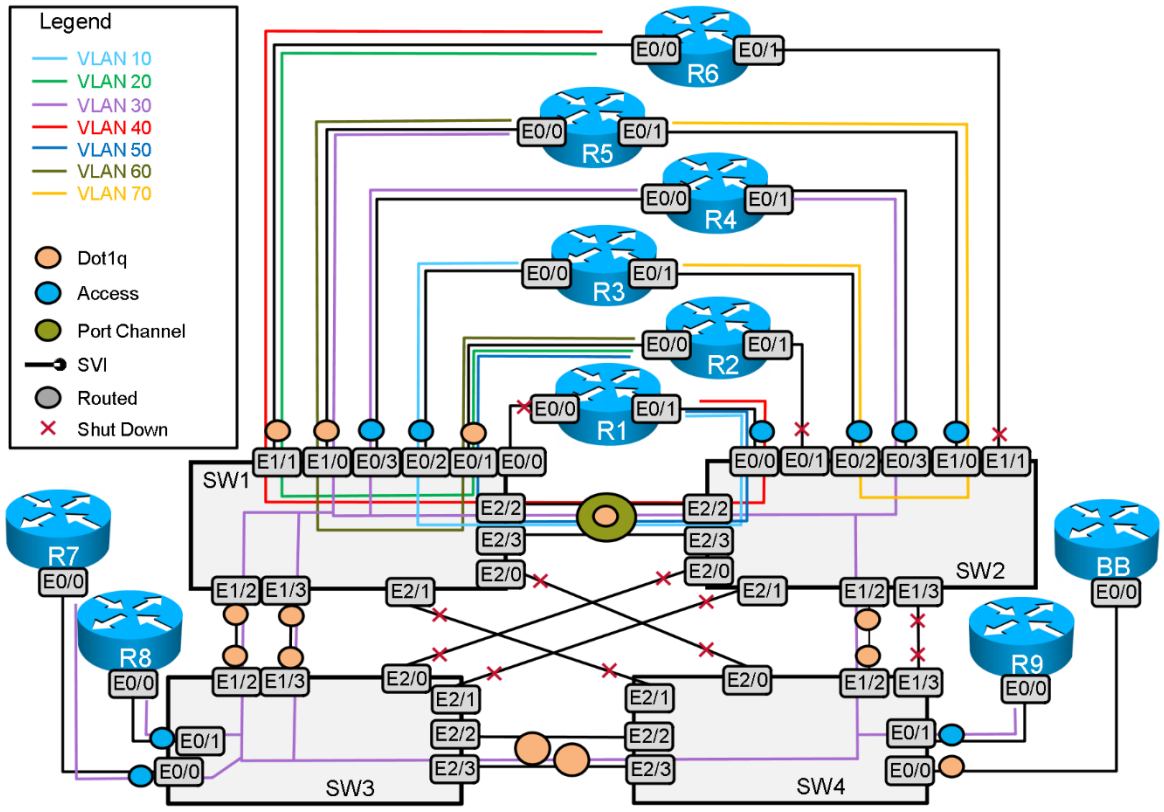

Carefully review the entire scenario. Closely examine the supplied diagram and any associated tables. Determine how you need to configure VTP, how to configure ports that are assigned as trunks, and how to configure ports that are assigned as static VLAN ports. Use the **switchport mode access** command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

Issue: Do not configure any switch for the VTP server mode.

Solution:

VLANs must be created on each switch manually. The VTP mode on the switches is transparent because server mode is not allowed. Therefore, switches will not be able to learn the VLAN information via VTP from each other.

SW1

```
vtp mode transparent
!
vlan 10
 name CA
!
vlan 20
 name VA
!
vlan 30
 name TX
!
vlan 40
 name MD
!
vlan 50
 name NY
!
vlan 60
 name MI
!
```

SW2

```
vtp mode transparent
!
vlan 10
 name CA
!
!
vlan 30
 name TX
!
vlan 40
 name MD
!
vlan 50
 name NY
!
!
vlan 70
 name DC
!
```

SW3

```
vtp mode transparent
!
vlan 30
 name TX
!
```

SW4


```
vtp mode transparent
!
vlan 30
 name TX
!
```

Verify VLAN names and check that all access ports are assigned to the proper VLANs.

```
SW1#show vlan brief | e (^[ 1] )|(^100[2-5])

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
10   CA                               active    Et0/2
20   VA                               active
30   TX                               active    Et0/3
40   MD                               active
50   NY                               active
60   MI                               active

SW2#show vlan brief | e (^[ 1] )|(^100[2-5])

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
10   CA                               active
30   TX                               active    Et0/3
40   MD                               active
50   NY                               active
70   DC                               active    Et0/2, Et1/0

SW3#show vlan brief | e (^[ 1] )|(^100[2-5])

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
30   TX                               active    Et0/0, Et0/1

SW4#show vlan brief | e (^[ 1] )|(^100[2-5])

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
30   TX                               active    Et0/1
```

Issue: Configure EtherChannel between SW1 and SW2, bundling ports 2/2 and 2/3 on both switches.
The EtherChannel should be built using the Cisco aggregation protocol with both switches actively participating in protocol negotiations.

Solution:

To configure EtherChannel, use the **channel-group** command and choose a mode. The **on**, **desirable**, and **auto** modes are Cisco Port Aggregation Protocol (PAgP) modes. PAgP is a Cisco proprietary protocol. Setting both sides to **desirable** unconditionally enables PAgP. It places an interface into an active negotiating state, in which the interface starts negotiations with other interfaces by sending PAgP packets:

```
SW1# conf t
SW1(config)# interface range E2/2 - 3
SW1(config-if-range)#channel-group 1 mode ?
  active     Enable LACP unconditionally
  auto       Enable PAgP only if a PAgP device is detected
  desirable  Enable PAgP unconditionally
  on         Enable Etherchannel only
  passive    Enable LACP only if a LACP device is detected
```

```
SW1(config-if-range)#channel-group 1 mode desirable
SW1(config-if-range)#end

SW2# conf t
SW2(config)# interface range E2/2 - 3
SW2(config-if-range)#channel-group 1 mode ?
  active     Enable LACP unconditionally
  auto       Enable PAgP only if a PAgP device is detected
  desirable  Enable PAgP unconditionally
  on         Enable Etherchannel only
  passive    Enable LACP only if a LACP device is detected

SW2(config-if-range)#channel-group 1 mode desirable
SW2(config-if-range)#end
```

The output below demonstrates that an EtherChannel was negotiated successfully:

```
SW1#show etherchannel summary
Flags:  D - down        P - bundled in port-channel
        I - stand-alone s - suspended
        H - Hot-standby (LACP only)
        R - Layer3      S - Layer2
        U - in use      N - not in use, no aggregation
        f - failed to allocate aggregator

        M - not in use, no aggregation due to minimum links not met
        m - not in use, port not aggregated due to minimum links not met
        u - unsuitable for bundling
        d - default port

        w - waiting to be aggregated
Number of channel-groups in use: 1
Number of aggregators:           1

Group  Port-channel  Protocol    Ports
------+-------------+-----------+-----------------------------------------------
-
1      Po1(SU)         PAgP      Et2/2(P)       Et2/3(P)
```

Issue: Use the 802.1Q protocol for all trunks whenever possible.

Solution:

The default Cisco virtual IOS Software on UNIX trunk encapsulation is negotiate mode, and it prefers Inter-Switch Link (ISL); the default trunk mode is desirable. If you did not statically configure the trunk ports, you are likely to find that ports connecting the switches have automatically negotiated ISL trunks with output similar to the following:

```
SW4#show interfaces trunk | i (^P)|(^$)|(2/2)

Port               Mode         Encapsulation  Status        Native vlan
Et2/2              desirable    n-isl          trunking      1

Port               Vlans allowed on trunk
Et2/2              1-4094
```

However, this lab requires that all trunks are configured with 802.1Q encapsulation whenever possible. In this scenario, all trunks are configured as 802.1Q trunks. It is a best practice to ensure that trunk encapsulation and mode are present on both sides of the trunk.

Configuration:
SW1

```
interface Port-channel1
```



```
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 10,30,40,50
!
interface Ethernet0/1
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 20,50,60
!
interface Ethernet1/0
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 30,60
!
interface Ethernet1/1
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 20,40
!
interface Ethernet1/2
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 30
!
interface Ethernet1/3
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 30
!
interface Ethernet2/2
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 10,30,40,50
!
interface Ethernet2/3
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 10,30,40,50
!
```

SW2

```
interface Port-channel1
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 10,30,40,50
interface Ethernet0/0
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 10,40,50
!
interface Ethernet1/2
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 30
!
interface Ethernet2/2
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 10,30,40,50
!
interface Ethernet2/3
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 10,30,40,50
!
```

SW3

```
interface Ethernet1/2
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 30
!
interface Ethernet1/3
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 30
!
interface Ethernet2/2
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 30
!
interface Ethernet2/3
```

```
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 30
!
```

SW4

```
interface Ethernet1/2
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 30
!
interface Ethernet2/2
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 30
!
interface Ethernet2/3
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 30
!
```

There are many commands that can help you confirm your trunk configuration. In the Mentor Guide engine, try the **show interfaces trunk** command on SW1 and SW2. Ensure that you can ping all of the interfaces within the same Ethernet subnet before moving on.

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

2. Internet Connectivity

Issue: R1 and R4 Internet connectivity

Solution:

This section requires basic Internet connectivity, which will later be used for GRE Tunnel VPN.

Note that this section explicitly allows use of static routes.

R1:

```
!
interface Serial1/1
 ip address 8.8.1.2 255.255.255.0
!
ip route 8.8.3.2 255.255.255.255 Serial1/1 8.8.1.1
!
```

R4:

```
interface Serial1/1
 ip address 8.8.3.2 255.255.255.0
!
ip route 8.8.1.2 255.255.255.255 Serial1/1 8.8.3.1
!
```

Note that the static host routes with the network mask 255.255.255.255 are configured on R1 and R4.

Issue: R1 and R2 Internet connectivity

Solution:

This section requires Internet connectivity, which will later be used for a GRE tunnel VPN. There are two differences between this section and the previous section:

- Internet connectivity needs to be done in a separate routing table.
- The use of static routing is not allowed, so ISP-advertised RIP routes need to be used instead.

Configuration example:

R1:

```
!
ip vrf Internet
!
interface Serial1/0
 ip vrf forwarding Internet
 ip address 8.8.0.2 255.255.255.0
!
router rip
 version 2
 no auto-summary
 !
 address-family ipv4 vrf Internet
  network 8.0.0.0
  no auto-summary
  version 2
 exit-address-family
!
```

R2:

```
!
ip vrf Internet
!

interface Serial1/0
 ip vrf forwarding Internet
 ip address 8.8.2.2 255.255.255.0
!
router rip
 version 2
 no auto-summary
 !
 address-family ipv4 vrf Internet
  network 8.0.0.0
  no auto-summary
  version 2
 exit-address-family
!
```

Verify routing and connectivity:

```
R1#show ip route vrf Internet

R*    0.0.0.0/0 [120/1] via 8.8.0.1, 00:00:09, Serial1/0
      8.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C        8.8.0.0/24 is directly connected, Serial1/0
L        8.8.0.2/32 is directly connected, Serial1/0

R1#ping vrf Internet 8.8.2.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 8.8.2.2, timeout is 2 seconds:
!!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 16/16/17 ms
```

Although the ISP advertises the default route, use of the default route is not restricted by the "Restriction and Goals" section, because this route is accepted in a separate routing table and is not shown with the **show ip route** command.

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

3. VPN Communications

Note	Re-examine VPN requirements and configuration when working on routing, IPv6, and multicast sections. The tunnel interfaces handle IPV4 unicast, IPv4 multicast and IPv6 packets differently. Verification of IPv4 unicast forwarding only (ping) is not sufficient verification to conclude that VPN is correctly configured and is suitable for all requirements of a scenario.

Issue: VPN connectivity between R1 and R2

Solution:

To fulfill requirements of this section, create tunnel interfaces using GRE (IP) encapsulation. Associate the tunnel with VRF "Internet" using the **tunnel vrf Internet** command. This command modifies VRF membership of GRE encapsulated packets. The packets within the tunnel remain in the global routing table (the VRF table can be assigned with the **ip vrf forwarding** command).

Configuration:

R1:

```
!
interface Tunnel12
 ip address 135.15.12.1 255.255.255.0
 tunnel source 8.8.0.2
 tunnel destination 8.8.2.2
 tunnel vrf Internet
!
```

R2:

```
!
interface Tunnel12
 ip address 135.15.12.2 255.255.255.0
 tunnel source 8.8.2.2
 tunnel destination 8.8.0.2
 tunnel vrf Internet
!
```

Verify the VPN connectivity. Here is an example on R1:

```
R1#ping 135.15.12.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 135.15.12.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 14/16/17 ms
```


Issue: VPN connectivity between R1 and R4

Solution:

Basic connectivity is established by creating tunnel interfaces using GRE (IP) encapsulation.

Configuration:

R1:

```
!
interface Tunnel14
 ip address 135.15.14.1 255.255.255.0
 tunnel source 8.8.1.2
 tunnel destination 8.8.3.2
!
```

R4:

```
!
interface Tunnel14
 ip address 135.15.14.4 255.255.255.0
 tunnel source 8.8.3.2
 tunnel destination 8.8.1.2
!
```

Issue: VPN connectivity between R4 and R7

Solution:

Configure the mGRE Tunnel10 interfaces on R4 and R7 according to the scenario requirements. The interface needs to be configured as mGRE to allow the future addition of routers onto this subnet. However, the use of dynamic Next Hop Resolution Protocol (NHRP) is restricted. Instead, create static next-hop mapping on all routers. Mapping needs to be created for IPv4 unicast and IPV4 multicast forwarding. If more routers are added in the future, static mapping will need to be created to accommodate new routers. The use of dynamic NHRP enables growth of the mGRE network.

Example configuration:

R4:

```
interface Tunnel10
 ip address 135.15.10.4 255.255.255.0
 ip nhrp map 135.15.10.7 135.15.20.7
 ip nhrp map multicast 135.15.20.7
 ip nhrp network-id 10
 tunnel source Ethernet0/0
 tunnel mode gre multipoint
 tunnel key 10
!
```

R7:

```
interface Tunnel10
 ip address 135.15.10.7 255.255.255.0
 ip nhrp map 135.15.10.4 135.15.20.4
 ip nhrp map multicast 135.15.20.4
 ip nhrp network-id 10
 tunnel source Ethernet0/0
 tunnel mode gre multipoint
 tunnel key 10
!
```

Verify crypto sessions on R1 and R4:

R1:

```
R1#show crypto session
Crypto session current status

Interface: Tunnel14
Session status: UP-ACTIVE
Peer: 8.8.3.2 port 500
  IKEv1 SA: local 8.8.1.2/500 remote 8.8.3.2/500 Active
  IPSEC FLOW: permit 47 host 8.8.1.2 host 8.8.3.2
        Active SAs: 4, origin: crypto map
```

R4:

```
R1#show crypto session
R4#show crypto session
Crypto session current status

Interface: Tunnel14
Session status: UP-ACTIVE
Peer: 8.8.1.2 port 500
  IKEv1 SA: local 8.8.3.2/500 remote 8.8.1.2/500 Active
  IPSEC FLOW: permit 47 host 8.8.3.2 host 8.8.1.2
        Active SAs: 4, origin: crypto map
```

Note that the crypto sessions are up and active on the Tunnel14 interfaces of R1 and R4.

Verify next-hop mapping and basic IPv4 unicast connectivity:

```
R4#show ip nhrp
135.15.10.7/32 via 135.15.10.7
   Tunnel10 created 00:05:42, never expire
   Type: static, Flags: used
   NBMA address: 135.15.20.7

R4#ping 135.15.10.7
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 135.15.10.7, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

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

4. VPN Security

After the tunnel interface is created, IPsec encryption needs to be applied.

IPsec can encrypt GRE packets using a crypto map or tunnel protection. Both methods specify that IPsec encryption is performed after GRE encapsulation is configured. When a crypto map is used, encryption is applied to the outbound physical interfaces for the GRE tunnel packets. When tunnel protection is used, encryption is configured on the GRE tunnel interface.

The requirements specify that protection cannot be associated with the physical interface, so the IPsec profile configuration method needs to be used.

Create the required ISAKMP and IPsec crypto configurations:

R1:

```
!
crypto isakmp policy 1
 encr 3des
 hash md5
 authentication pre-share
 group 2
!
crypto isakmp key R1-R4-secret address 8.8.3.2
!
crypto ipsec transform-set ts_01 ah-sha-hmac esp-3des
 mode transport
!
crypto ipsec profile ipsec_prof
 set transform-set ts_01
!
```

R4:

```
crypto isakmp policy 1
 encr 3des
 hash md5
 authentication pre-share
 group 2
!
crypto isakmp key R1-R4-secret address 8.8.1.2
!
crypto ipsec transform-set ts_01 ah-sha-hmac esp-3des
 mode transport
!
crypto ipsec profile ipsec_prof
 set transform-set ts_01
!
```

Attach the IPsec profile to Tunnel interface:

R1:

```
!
interface Tunnel14
 tunnel protection ipsec profile ipsec_prof
!
```

R4:

```
!
interface Tunnel14
 tunnel protection ipsec profile ipsec_prof
!
```

Verify connectivity:

```
R1#ping 135.15.12.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 135.15.12.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 14/16/17 ms
```

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

5. IPv4 OSPF

Issue: Configure OSPF Area 0 between R4 and R7 on the network 135.15.10.0/24.
On Area 0, use the OSPF network type that elects a designated router (DR) or backup designated router (BDR) and does not use the unicast packet exchange.
R7 should be configured not to become the DR.

Solution:

Configuration:

R4:

```
!
interface Tunnel10
 ip ospf network broadcast
!
router ospf 1
 network 135.15.10.0 0.0.0.255 area 0
!
```

R7:

```
!
interface Tunnel10
 ip ospf network broadcast
 ip ospf priority 0
!
router ospf 1
 network 135.15.10.0 0.0.0.255 area 0
!
```

Issue: Configure OSPF Area 10 on links 135.15.13.0/24, 135.15.14.0/24, and 135.15.34.0/24.
The point-to-point OSPF network type should be used on 135.15.13.0/24 and 135.15.14.0/24.
The broadcast OSPF network type should be used on 135.15.34.0/24, with R4 as the DR.

Solution:

Configuration:

R1:

```
!
interface Ethernet0/1.10
 ip address 135.15.13.1 255.255.255.0
 ip ospf network point-to-point
!
interface Tunnel14
 ip ospf network point-to-point
!
router ospf 1
 network 135.15.13.0 0.0.0.255 area 10
 network 135.15.14.0 0.0.0.255 area 10
!
```

R3:

```
!
router ospf 1
 network 135.15.13.0 0.0.0.255 area 10
 network 135.15.34.0 0.0.0.255 area 10
!
interface Ethernet0/0
 ip address 135.15.13.3 255.255.255.0
```



```
 ip ospf network point-to-point
!
interface Serial1/0
 ip address 135.15.34.3 255.255.255.0
 ip ospf network broadcast
 ip ospf priority 0
!
router ospf 1
 network 135.15.13.0 0.0.0.255 area 10
 network 135.15.34.0 0.0.0.255 area 10
!
```

R4:

```
!
interface Tunnel14
 ip ospf network point-to-point
!
interface Serial1/0
 ip address 135.15.34.4 255.255.255.0
 ip ospf network broadcast
!
router ospf 1
 network 135.15.14.0 0.0.0.255 area 10
 network 135.15.34.0 0.0.0.255 area 10
!
```

Issue: Originate OSPF traffic from R1 on subnet 135.15.12.0/24.

Solution:

You should configure a nonzero OSPF priority as well as a neighbor statement on router R1, and OSPF priority 0 on R2, so that R2 will never initiate hello packets on the 135.15.12.0/24 subnet.

Configure OSPF Area 20 and a neighbor statement on R1:

```
router ospf 1
 network 135.15.12.0 0.0.0.255 area 20
 neighbor 135.15.12.2
interface Tunnel12
 ip address 135.15.12.1 255.255.255.0
 ip ospf network non-broadcast
```

Configure OSPF Area 20 and OSPF priority 0 on R2:

```
interface Tunnel12
 ip address 135.15.12.2 255.255.255.0
 ip ospf network non-broadcast
 ip ospf priority 0
!
router ospf 1
 network 135.15.12.0 0.0.0.255 area 20
```

Issue: OSPF Area 20 is not directly connected to OSPF Area 0*.*

Solution:

Configure a virtual link through OSPF Area 10.

R1:

```
!
router ospf 1
```

```
 area 10 virtual-link 135.15.104.1
!
```

R4:

```
!
router ospf 1
 area 10 virtual-link 135.15.101.1
!
```

Verify the virtual link:

```
R1#show ip ospf virtual-links

Virtual Link OSPF_VL0 to router 135.15.104.1 is up
  Run as demand circuit
  DoNotAge LSA allowed.
  Transit area 10, via interface Ethernet0/1.10
 Topology-MTID    Cost    Disabled     Shutdown      Topology Name
        0          74        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 00:00:07
    Adjacency State FULL (Hello suppressed)
    Index 1/3, 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
```

Issue: The adjacency on the 135.15.14.0/24 link should be authenticated with a cleartext password.
The adjacency on the 135.15.34.0/24 link should be authenticated with MD5.

Solution:

Use the interface configuration command **ip ospf authentication-key** on the interfaces connected to subnet 135.15.14.0/24. This will fulfill the cleartext password configuration requirement.

R1

```
interface Tunnel14
 ip address 135.15.14.1 255.255.255.0
 ip ospf authentication
 ip ospf authentication-key clear-do
```

R4

```
interface Tunnel14
 ip address 135.15.14.4 255.255.255.0
 ip ospf authentication
 ip ospf authentication-key clear-do
```

Use the interface command **ip ospf message-digest-key 1 md5** on the interfaces connected to subnet 135.15.34.0/24.

R3

```
interface Serial1/0
 ip address 135.15.34.3 255.255.255.0
 ip ospf authentication message-digest
 ip ospf message-digest-key 1 md5 md5-test
```

R4

```
interface Serial1/0
 ip address 135.15.34.4 255.255.255.0
 ip ospf authentication message-digest
 ip ospf message-digest-key 1 md5 md5-test
```


Issue the **show ip ospf interface** command on R4:

```
R4#show ip ospf interface
Serial1/0 is up, line protocol is up
  Internet Address 135.15.34.4/24, Area 10, Attached via Network Statement
  Process ID 1, Router ID 135.15.104.1, Network Type BROADCAST, Cost: 64
   [skipped]
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 135.15.103.1
  Suppress hello for 0 neighbor(s)
  Message digest authentication enabled
    Youngest key id is 1
Tunnel14 is up, line protocol is up
  Internet Address 135.15.14.4/24, Area 10, Attached via Network Statement
  Process ID 1, Router ID 135.15.104.1, Network Type POINT_TO_POINT, Cost: 1000
  [skipped]
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 135.15.101.1
  Suppress hello for 0 neighbor(s)
  Simple password authentication enabled
```

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

6. IPv4 EIGRP

Issue: Configure Enhanced Interior Gateway Protocol (EIGRP) AS 30 on R4 and R8 interfaces connected to VLAN 30.

Solution:

Configure the **router eigrp 30** command and advertise network 135.10.20.0 with the 0.0.0.255 wildcard mask on routers R4 and R8.

R4

```
router eigrp 30
 network 135.15.20.0 0.0.0.255
!
```

R8

```
router eigrp 30
 network 135.15.20.0 0.0.0.255
!
```

Issue: Configure EIGRP AS 162 on the interfaces connected to VLANs 20 and 40 between the routers R1, R6, and R2, as well as on the Loopback 106 interface of the router R6.

Solution:

Configure **router eigrp 162** and advertise specified networks with the wildcard.

R1

```
router eigrp 162
 network 135.15.16.0 0.0.0.255
!
```

R2
```
router eigrp 162
 network 135.15.26.0 0.0.0.255
!
```

R6
```
router eigrp 162
 network 135.15.16.0 0.0.0.255
 network 135.15.26.0 0.0.0.255
 network 135.15.106.0 0.0.0.255
!
```

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

7. IPv4 RIP

Issue: Configure RIP version 2 (RIPv2).

Solution:

There are two options for configuring RIPv2. Either specify the version on the RIP-enabled interfaces of the RIP-speaking router or specify "version 2" in the RIP router configuration process.

Use the **router rip passive interface** command under the RIP routing process to make RIP advertise updates only to certain interfaces.

Configuration:

R2
```
router rip
 version 2
 passive-interface default
 no passive-interface Ethernet0/0.60
 no passive-interface Serial1/1
 network 135.15.0.0
 no auto-summary
!
```

R3
```
router rip
 version 2
 passive-interface default
 no passive-interface Ethernet0/1
 no passive-interface Serial1/1
 network 135.15.0.0
 no auto-summary
!
```

R5
```
router rip
 version 2
 passive-interface default
 no passive-interface Ethernet0/0.30
 no passive-interface Ethernet0/0.60
```



```
 no passive-interface Ethernet0/1
 network 135.15.0.0
 no auto-summary
!
```

R9

```
router rip
 version 2
 passive-interface default
 no passive-interface Ethernet0/0
 network 135.15.0.0
 no auto-summary
!
```

Verification:

```
R2#show ip protocols | s "rip
Routing Protocol is "rip"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Sending updates every 30 seconds, next due in 26 seconds
  Invalid after 180 seconds, hold down 0, flushed after 240
  Default redistribution metric is 1
  Redistributing: rip
  Default version control: send version 2, receive version 2
    Interface             Send  Recv  Triggered RIP  Key-chain
    Ethernet0/0.60        2     2
    Serial1/1             2     2
  Automatic network summarization is not in effect
  Maximum path: 4
  Routing for Networks:
     135.15.0.0
    135.15.0.0
  Passive Interface(s):
    Ethernet0/0
    Ethernet0/0.20
    Ethernet0/0.50
    Ethernet0/1
    Ethernet0/2
    Ethernet0/3
    Serial1/2
    Serial1/3
    Loopback2
    Loopback102
    RG-AR-IF-INPUT1
    Tunnel0
    Tunnel12
    VoIP-Null0
  Routing Information Sources:
    Gateway         Distance      Last Update
    135.15.25.5           90      00:00:19
    135.15.23.3           90      00:00:04
  Distance: (default is 120)
```

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

8. 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 ensure that no routing loops exist and, whenever possible, ensure that no suboptimal paths are selected.

Three IGPs are configured in this scenario. Two of these IGPs are classified as core routing protocols (OSPF, RIP), and one is classified as an edge routing protocol (EIGRP, including AS 162 and AS 30). The core routing protocols provide transit services to the edge protocols. The core routing protocols might also provide a level of redundancy to the edge protocols. Four redistribution points are selected to perform redistribution: R1, R2, R3, and R4.

R1 connects EIGRP 162 and OSPF. R2 connects EIGRP 162, OSPF, and RIP. Since EIGRP is an edge protocol in this scenario, EIGRP will not provide transit services to the rest of the internetwork. R1 and R2 are almost equal from the perspective of a redistribution process. The difference is that R1 will need to redistribute the connected network 111.1.1.1 into all protocols to support the multicast task. When you perform a **redistribute connected** operation with an accompanying route map, ensure that other prefixes are not blocked by the route map's access list. The route map that redistributes connected networks is checked first and will block all networks that are not permitted.

The administrative distance of native RIP prefixes is changed to 90 on R2. This value is less than the OSPF default administrative distance (110). The adjustment of the RIP administrative distance prevents the override of these native prefixes by OSPF in the routing table of R2.

R1 redistributes EIGRP prefixes from EIGRP AS 162 into OSPF. R2 redistributes EIGRP prefixes from EIGRP AS 162 into RIP and OSPF. EIGRP AS 162 receives all prefixes except its native prefixes from both RIP and OSPF; RIP redistributes OSPF-only prefixes from OSPF; and OSPF receives RIP-only prefixes from RIP. This type of filtering reduces the level of redundancy in the network, but also prevents the formation of routing loops, which is more important to the stability of the network.

R3 has OSPF and RIP redistributing mutually. RIP will receive all prefixes except its native prefixes from OSPF; OSPF will receive all prefixes except its native prefixes from RIP. There is also a need to redistribute the connected network 111.1.1.1 into both RIP and OSPF to support the multicast configuration task. Once again, when you use the **redistribute connected** command with a route map, ensure that other prefixes are not blocked by the route map's access list. A route map that redistributes connected networks is checked first and will block all networks that are not permitted. The administrative distance of native RIP prefixes is changed to 80 on R3, which is less than the OSPF default administrative distance (110) to prevent OSPF from overriding these prefixes in the R3 routing table.

R4 has OSPF and EIGRP redistributed mutually. Since EIGRP is an edge protocol in this scenario, EIGRP will not provide transit services to the rest of the internetwork. There is a need to redistribute the connected network 111.1.1.1 into both EIGRP and OSPF to support the multicast configuration task. Remember to watch for possible unwanted redistribution filters when using route maps to redistribute connected networks.

Loopbacks are advertised in different manners:

- Administrative Loopbacks 102.1, 103.1, 105.1, and 120.1 do not need to be redistributed as connected; the RIP process will inject them because they are subnets of a classful network defined under router RIP.

- Administrative Loopback 106.1 is part of EIGRP and is defined under router EIGRP.
- Administrative Loopbacks 101.1 and 104.1 are redistributed into OSPF and EIGRP as connected.
- Administrative Loopbacks 107.1 is redistributed into OSPF as connected.
- Administrative Loopbacks 110.1 is redistributed into EIGRP as connected.

Redistribution Table

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

IPv4 IGP Redistribution

Redist. Point	Into RIP		Into OSPF		Into EIGRP 30		Into EIGRP 162	
	Permit	Deny	Permit	Deny	Permit	Deny	Permit	Deny
R1			EIGRP AS 162 native Connected	(*)[1]			(*)[2]	EIGRP AS 162 only
R2	EIGRP 162 OSPF[3]	(*)[1]	Connected RIP only	(*)[1]			(*)[2]	EIGRP AS 162 only
R3	Connected OSPF[3]		Connected RIP only	(*)[1]				
R4			Connected EIGRP[3]		Connected OSPF[3]			

(*)[1] Others explicitly denied (*)[2] Others implicitly allowed [3] No filtering

R1:

```
!
router eigrp 162
 default-metric 1000 100 255 3 1500
 redistribute connected route-map Connected-2-EIGRP162
 redistribute ospf 1
!
router ospf 1
 redistribute connected subnets route-map Connected-2-OSPF
 redistribute eigrp 162 subnets route-map EIGRP162-2-OSPF
!
ip access-list standard Connected-2-EIGRP162
 permit 111.1.1.1
 permit 135.15.12.0
 permit 135.15.13.0
 permit 135.15.14.0
 permit 135.15.21.0
 permit 135.15.101.0
ip access-list standard Connected-2-OSPF
 permit 111.1.1.1
 permit 135.15.16.0
 permit 135.15.21.0
 permit 135.15.101.0
ip access-list standard EIGRP162-2-OSPF
```

```
 permit 135.15.26.0
 permit 135.15.106.0
!
route-map Connected-2-EIGRP162 permit 10
 match ip address Connected-2-EIGRP162
!
route-map EIGRP162-2-OSPF permit 10
 match ip address EIGRP162-2-OSPF
!
route-map Connected-2-OSPF permit 10
 match ip address Connected-2-OSPF
!
```

R2:

```
!
router eigrp 162
 default-metric 1000 100 255 3 1500
 redistribute ospf 1
 redistribute rip route-map RIP-2-EIGRP162
!
router ospf 1
 redistribute eigrp 162 subnets route-map EIGRP162-2-OSPF
 redistribute rip subnets route-map RIP-2-OSPF
!
router rip
 redistribute eigrp 162 route-map EIGRP162-2-RIP
 redistribute ospf 1
 distance 90 0.0.0.0 255.255.255.255 RIP-distance
!
ip access-list standard EIGRP162-2-OSPF
 permit 135.15.26.0
 permit 135.15.16.0
 permit 135.15.106.0
ip access-list standard EIGRP162-2-RIP
 permit 135.15.26.0
 permit 135.15.16.0
 permit 135.15.106.0
ip access-list standard RIP-2-EIGRP162
 deny   135.15.26.0
 deny   135.15.16.0
 deny   135.15.106.0
 permit any
ip access-list standard RIP-2-OSPF
 permit 135.15.25.0
 permit 135.15.20.0
 permit 135.15.21.0
 permit 135.15.23.0
 permit 135.15.35.0
 permit 135.15.105.0
 permit 135.15.102.0
 permit 135.15.103.0
 permit 135.15.120.0
ip access-list standard RIP-distance
 permit 135.15.20.0
 permit 135.15.35.0
 permit 135.15.105.0
 permit 135.15.110.0
 permit 135.15.103.0
 permit 135.15.120.0
!
route-map EIGRP162-2-RIP permit 10
 match ip address EIGRP162-2-RIP
!
route-map RIP-2-EIGRP162 permit 10
 match ip address RIP-2-EIGRP162
!
route-map EIGRP162-2-OSPF permit 10
 match ip address EIGRP162-2-OSPF
```



```
!
route-map RIP-2-OSPF permit 10
 match ip address RIP-2-OSPF
!
```

R3:

```
!
router ospf 1
 redistribute connected subnets route-map Connected-2-OSPF
 redistribute rip subnets route-map RIP-2-OSPF
!
router rip
 redistribute connected route-map Connected-2-RIP
 redistribute ospf 1
 distance 80 0.0.0.0 255.255.255.255 RIP-distance
!
ip access-list standard Connected-2-OSPF
 permit 111.1.1.1
 permit 135.15.23.0
 permit 135.15.35.0
 permit 135.15.103.0
ip access-list standard Connected-2-RIP
 permit 111.1.1.1
ip access-list standard RIP-2-OSPF
 permit 135.15.25.0
 permit 135.15.20.0
 permit 135.15.21.0
 permit 135.15.23.0
 permit 135.15.35.0
 permit 135.15.105.0
 permit 135.15.102.0
 permit 135.15.103.0
 permit 135.15.120.0
ip access-list standard RIP-distance
 permit 135.15.25.0
 permit 135.15.26.0
 permit 135.15.20.0
 permit 135.15.21.0
 permit 135.15.105.0
 permit 135.15.102.0
 permit 135.15.120.0
!
route-map Connected-2-RIP permit 10
 match ip address Connected-2-RIP
!
route-map Connected-2-OSPF permit 10
 match ip address Connected-2-OSPF
!
route-map RIP-2-OSPF permit 10
 match ip address RIP-2-OSPF
!
```

R4:

```
router eigrp 30
 default-metric 10000 10 255 1 1500
 redistribute connected route-map Connected-2-EIGRP
 redistribute ospf 1
!
router ospf 1
 redistribute connected subnets route-map Connected-2-OSPF
 redistribute eigrp 30 subnets
!
ip access-list standard Connected-2-EIGRP
 permit 111.1.1.1
 permit 135.15.10.0
 permit 135.15.14.0
 permit 135.15.34.0
 permit 135.15.104.0
ip access-list standard Connected-2-OSPF
```

```
 permit 111.1.1.1
 permit 135.15.20.0
 permit 135.15.104.0
route-map Connected-2-EIGRP permit 10
 match ip address Connected-2-EIGRP
!
route-map Connected-2-OSPF permit 10
 match ip address Connected-2-OSPF
!
```

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

9. Border Gateway Protocol

BGP Diagram

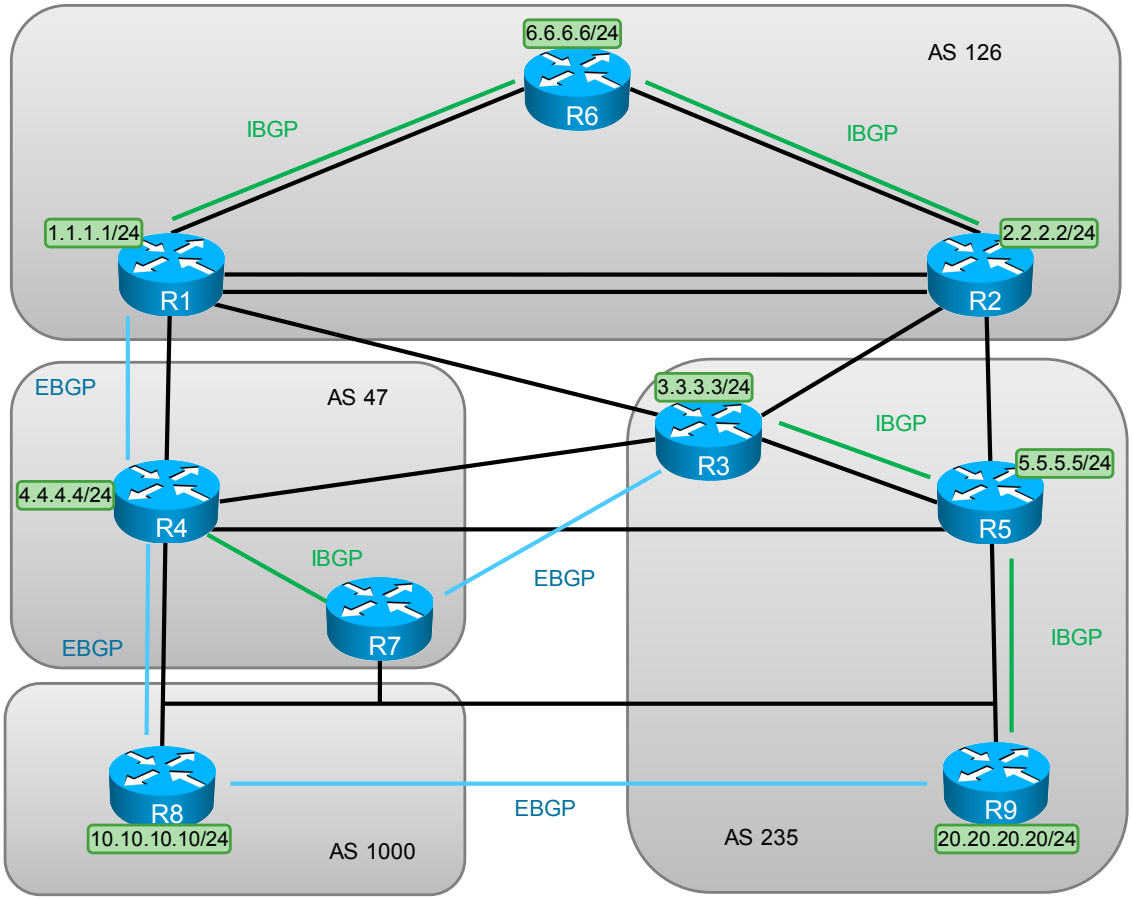

Issue: Peer relationship is not specified within AS 126 and AS 235.

Solution:

To fulfill the IBGP neighbor relationship formation requirements of both AS 126 and 235, you can choose from the following three configuration options: configure a full mesh of IBGP

speakers within the respective AS; configure a route reflector within the respective AS; or configure a confederation within the respective AS.

The answer key uses the route reflector configuration approach: R5 and R6 are configured as a route reflector in AS 235 and 126 respectively.

R1:

```
router bgp 126
 neighbor 135.15.106.1 remote-as 126
 neighbor 135.15.106.1 update-source Loopback101
!
```

R2:

```
router bgp 126
 neighbor 135.15.26.6 remote-as 126
!
```

R6:

```
router bgp 126
 neighbor 135.15.26.2 remote-as 126
 neighbor 135.15.26.2 route-reflector-client
 neighbor 135.15.101.1 remote-as 126
 neighbor 135.15.101.1 update-source Loopback106
 neighbor 135.15.101.1 route-reflector-client
!
```

R3:

```
router bgp 235
 neighbor 135.15.105.1 remote-as 235
 neighbor 135.15.105.1 update-source Loopback103
!
```

R5:

```
router bgp 235
 neighbor 135.15.103.1 remote-as 235
 neighbor 135.15.103.1 update-source Loopback105
 neighbor 135.15.103.1 route-reflector-client
 neighbor 135.15.120.1 remote-as 235
 neighbor 135.15.120.1 update-source Loopback105
 neighbor 135.15.120.1 route-reflector-client
!
```

R9:

```
router bgp 235
 neighbor 135.15.105.1 remote-as 235
 neighbor 135.15.105.1 update-source Loopback120
!
```

Issue: Configure R5 to ensure that R3 has a next hop 135.15.35.5 for subnet 20.20.20.0/24.

Solution:

Configure a route map outbound from R5 to R3 to set the next hop for subnet 20.20.20.0 to 135.15.35.5.

```
ip access-list standard NET20-NH
 permit 20.20.20.0 0.0.0.255
!
route-map NET20-NH permit 10
 match ip address NET20-NH
 set ip next-hop 135.15.35.5
!
route-map NET20-NH permit 20
!
router bgp 235
 neighbor 135.15.103.1 route-map NET20-NH out
!
```

Issue: Configure R7 to ensure that R3 has a next hop 135.15.106.1 for subnet 1.1.1.0/24, 2.2.2.0/24, 6.6.6.0/24.

Solution:

Configure a route map outbound from R7 to R3 to set the next hop.

```
ip access-list standard NET1-NH
 permit 1.1.1.0 0.0.0.255
 permit 2.2.2.0 0.0.0.255
 permit 6.6.6.0 0.0.0.255
!
route-map NET1-NH permit 10
 match ip address NET1-NH
 set ip next-hop 135.15.106.1
!
route-map NET1-NH permit 20
!
router bgp 47
 neighbor 135.15.103.1  route-map NET1-NH out
!
```

Issue: Configure R4 to ensure that R7 has a next hop 135.15.20.4 for subnet 1.1.1.0/24, 2.2.2.0/24, 6.6.6.0/24.

Solution:

Configure a route map outbound from R4 to R7 to set the next hop.

R4

```
ip access-list standard NET1-NH
 permit 1.1.1.0 0.0.0.255
 permit 2.2.2.0 0.0.0.255
 permit 6.6.6.0 0.0.0.255
!
route-map NET1-NH permit 10
 match ip address NET1-NH
 set ip next-hop 135.15.20.4
!
route-map NET1-NH permit 20
!
router bgp 47
 neighbor 135.15.107.1 route-map NET1-NH out
```


!

Note	Before this section is completed, you can observe routing loops while trying to ping BGP subnets. Attempt to ping addresses 1.1.1.1, 2.2.2.2, 6.6.6.6, and 20.20.20.20 from R3 and R7 and attempt to find and explain the loops.

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

10. Control Traffic

Issue: If traffic is originated from the 135.15.103.0/24 subnet and destined to the 135.15.106.0/24 subnet, it must be forwarded to R2 from R3.

Solution:

Local policy routing can be used on R3 to forward traffic destined to 135.15.106.0/24 to R2.

R3

```
ip local policy route-map Traffic-control
!
ip access-list extended Traffic-control
 permit ip 135.15.103.0 0.0.0.255 135.15.106.0 0.0.0.255
!
route-map Traffic-control permit 10
 match ip address Traffic-control
 set ip next-hop 135.15.23.2
!
```

Issue: This traffic flow must traverse router R2 twice and R1 once in one direction. This traffic flow must be symmetric.

Solution:

This traffic must enter R2 twice in the direction from R3 to R6 and must enter R2 twice in the direction from R6 to R3. The set of traversing incoming and outgoing interfaces should be consistent along the forwarding path

Issue: Do not make forwarding decisions based on the precise source and destination match on R1.

Solution:

Considering all the previous issues, a solution that optimally meets the configuration requirements is needed.

Traffic forwarded from 135.15.103.0/24 to 135.15.106.0/24 will enter R2 on the interface connected to subnet 135.15.23.0/24. You cannot forward it to R6 because the requirement says that the traffic must enter R2 twice. You need to forward traffic to another intermediate router. R1 is a good candidate. To forward traffic to R1 via the interface on the 135.15.12.0/24 network, you can use the policy routing based on the prefix match on the interface on 135.15.23.2/24. On R1, you can configure policy routing as well, but you cannot use the match criteria based on a packet's source and destination address. To determine which criteria you can use in this task, read the QoS section. The IP traffic flow terminated between 135.15.103.0/24 and 135.15.106.0/24 will be marked with the IP precedence of 2; therefore, the IP precedence can be used as match criteria for the policy routing on R1. R1 will forward traffic back to R2. R1 has two options to forward traffic back to R2: forward traffic back out the Frame Relay interface connected to 135.15.12.0/24, or forward traffic via the VLAN 50 interface. The VLAN 50 interface will be used in this answer key. The following diagram illustrates the implementation plan:

Traffic Control Diagram

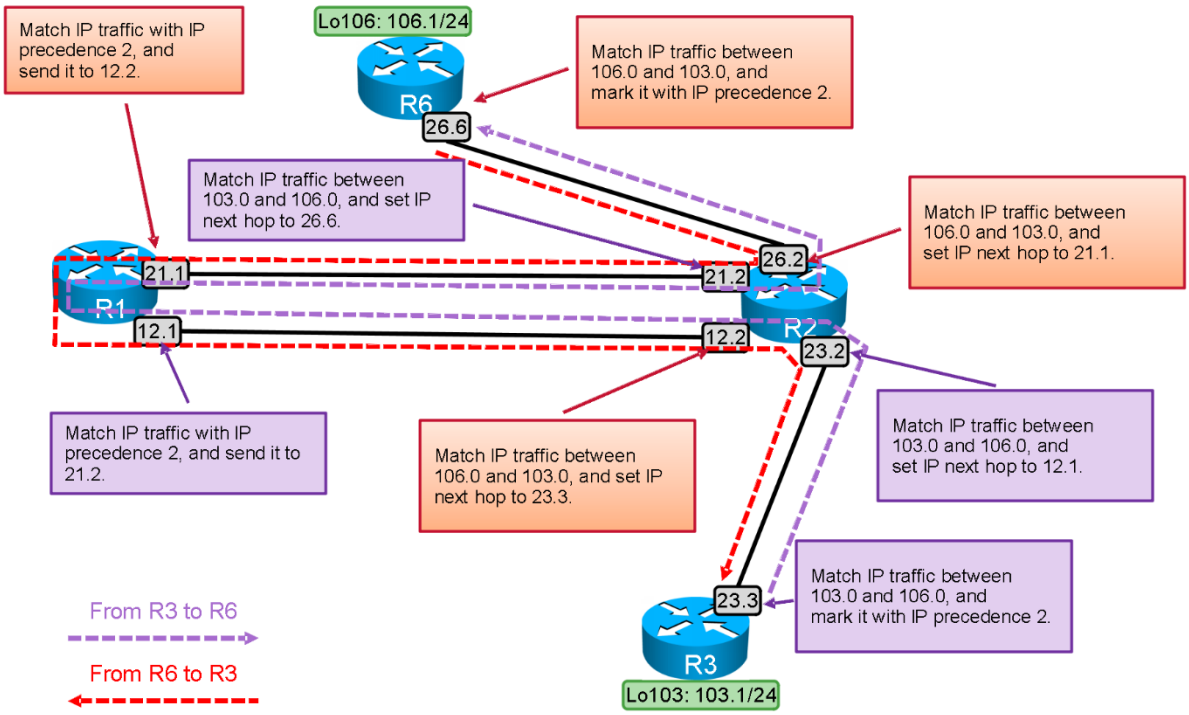

- Traffic (135.15.103.0/24 → 135.15.106.0/24) is originated on R3 and is forwarded to R2 via 135.15.23.2 as the next hop. This task is accomplished by using the **local policy** routing configuration command on R3. Traffic between networks 135.15.103.0/24 and 135.15.106.0 will be marked with precedence 2, which is done using QoS and without any route maps involved.

```
class-map match-all QOS
  match access-group name QOS
!
policy-map QOS
  class QOS
   set ip precedence 2
!
interface Serial1/1
 service-policy output QOS
!
ip local policy route-map Traffic-control
!
ip access-list extended Traffic-control
 permit ip 135.15.103.0 0.0.0.255 135.15.106.0 0.0.0.255
!
ip access-list extended QOS
 permit ip 135.15.103.0 0.0.0.255 135.15.106.0 0.0.0.255

route-map Traffic-control permit 10
```

```
 match ip address Traffic-control
 set ip next-hop 135.15.23.2
!
```

- Traffic from 135.15.103.0/24 to 135.15.106.0/24 enters R2 for the first time and is forwarded to R1 via 135.15.12.1 as the next hop. This task is accomplished by an interface policy routing configuration on R2 based on the prefix match criteria.

```
interface Serial1/1
 ip policy route-map Traffic-control_S0
   !
ip access-list extended Traffic-control_R3-2-R6
 permit ip 135.15.103.0 0.0.0.255 135.15.106.0 0.0.0.255
!
route-map Traffic-control_S0 permit 10
 match ip address Traffic-control_R3-2-R6
 set ip next-hop 135.15.12.1
!
```

- Traffic enters R1 and is forwarded back to R2 via 135.15.21.2 as the next hop on the VLAN 50 link. This task is accomplished by the interface policy routing configuration on R1 based on the IP precedence match criteria.

```
interface Tunnel12
 ip policy route-map Traffic-control-S0/0/0.12
!
ip access-list extended Traffic-control
 permit ip any any precedence immediate
!
route-map Traffic-control-Tu12 permit 10
 match ip address Traffic-control
 set ip next-hop 135.15.21.2
```

- Traffic from 135.15.103.0/24 to 135.15.106.0/24 enters R2 for the second time and is forwarded to R6 via 135.15.26.6 as the next hop. This task is accomplished by an interface policy routing configuration on R2 based on the prefix match criteria.

```
interface Ethernet0/0.50
 ip policy route-map Traffic-control_VLAN50
!
ip access-list extended Traffic-control_R3-2-R6
 permit ip 135.15.103.0 0.0.0.255 135.15.106.0 0.0.0.255
!
route-map Traffic-control_VLAN50 permit 10
 match ip address Traffic-control_R3-2-R6
 set ip next-hop 135.15.26.6
!
```

- The return traffic (135.15.106.0/24 → 135.15.103.0/24) is policy-routed similarly to meet the symmetric traffic requirement. See the diagram and the Mentor Guide for the configuration details.

Follow these steps to verify your configuration:

Generate IP traffic from R3 to R6:

```
ping 135.15.106.1 sour 135.15.103.1  rep 10
```

On R3, verify local policy and traffic marking:


```
R3#show ip local policy
Local policy routing is enabled, using route map Traffic-control
route-map Traffic-control, permit, sequence 10
  Match clauses:
    ip address (access-lists): Traffic-control
  Set clauses:
    ip next-hop 135.15.23.2
  Policy routing matches: 10 packets, 1000 bytes

R3#show policy-map in s1/1
Serial1/1

  Service-policy output: QOS

    Class-map: QOS (match-all)
      10 packets, 1040 bytes
      5 minute offered rate 0000 bps, drop rate 0000 bps
      Match: access-group name QOS
      QoS Set
        precedence 2
          Packets marked 10

    Class-map: class-default (match-any)
      783 packets, 89804 bytes
      5 minute offered rate 0000 bps, drop rate 0000 bps
      Match: any
```

On R2, verify the routing policy:

```
R2#show ip policy
Interface      Route map
Ethernet0/0.20 Traffic-control_VLAN20
Ethernet0/0.50 Traffic-control_VLAN50
Serial1/1      Traffic-control_S0
Tunnel12       Traffic-control_Tu12

R2#show route-map Traffic-control_S0
route-map Traffic-control_S0, permit, sequence 10
  Match clauses:
    ip address (access-lists): Traffic-control_R3-2-R6
  Set clauses:
    ip next-hop 135.15.12.1
  Policy routing matches: 10 packets, 1040 bytes

R2#show route-map Traffic-control_VLAN50
route-map Traffic-control_VLAN50, permit, sequence 10
  Match clauses:
    ip address (access-lists): Traffic-control_R3-2-R6
  Set clauses:
    ip next-hop 135.15.26.6
  Policy routing matches: 10 packets, 1180 bytes

R2#show route-map Traffic-control_VLAN20
route-map Traffic-control_VLAN20, permit, sequence 10
  Match clauses:
    ip address (access-lists): Traffic-control_R6-2-R3
  Set clauses:
    ip next-hop 135.15.21.1
  Policy routing matches: 10 packets, 1180 bytes

R2# show route-map Traffic-control_Tu12
route-map Traffic-control_Tu12, permit, sequence 10
  Match clauses:
    ip address (access-lists): Traffic-control_R6-2-R3
  Set clauses:
    ip next-hop 135.15.23.3
```

```
  Policy routing matches: 10 packets, 1000 bytes
```

R1 should match on the precedence criteria:

```
R1#show ip policy
Interface      Route map
Ethernet0/1.50 Traffic-control-VLAN50
Tunnel12       Traffic-control-Tu12

R1#show route-map Traffic-control-VLAN50
route-map Traffic-control-VLAN50, permit, sequence 10
  Match clauses:
    ip address (access-lists): Traffic-control
  Set clauses:
    ip next-hop 135.15.12.2
  Policy routing matches: 10 packets, 1180 bytes

R1#show route-map Traffic-control-Tu12
route-map Traffic-control-Tu12, permit, sequence 10
  Match clauses:
    ip address (access-lists): Traffic-control
  Set clauses:
    ip next-hop 135.15.21.2
  Policy routing matches: 10 packets, 1000 bytes


R1#show access-list Traffic-control
Extended IP access list Traffic-control
    10 permit ip any any precedence immediate (20 matches)
    20 permit ip 135.15.103.0 0.0.0.255 135.15.106.0 0.0.0.255
    30 permit ip 135.15.106.0 0.0.0.255 135.15.103.0 0.0.0.255
```

Verify the traffic marking on R6:

```
R6#show interface E0/0.20 rate
Ethernet0/0.20 VLAN20
  Output
    matches: access-group 101
      params:  100000000 bps, 18750000 limit, 25000000 extended limit
      conformed 10 packets, 1180 bytes; action: set-prec-transmit 2
      exceeded 0 packets, 0 bytes; action: set-prec-transmit 2
      last packet: 490463ms ago, current burst: 0 bytes
      last cleared 00:37:00 ago, conformed 4 bps, exceeded 0 bps
```

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

11. IPv4 Connectivity Verification

Below is a Tool Command Language (Tcl) script to test universal reachability. To use the script, enter the **tclsh** command in privileged mode and paste this script. To kill failing pings, press **Ctrl-Shift** while pressing **6** twice. When you are done, enter the **tclquit** command to leave Tcl mode.

Note	This script includes all IPv4 addresses, including addresses listed in the following section of this scenario (that is, BGP addresses).

```
tclsh
foreach address {
1.1.1.1
111.1.1.1
135.15.101.1
```



```
135.15.12.1
135.15.14.1
135.15.13.1
135.15.16.1
135.15.21.1
2.2.2.2
135.15.102.1
135.15.12.2
135.15.26.2
135.15.21.2
135.15.25.2
135.15.23.2
3.3.3.3
135.15.103.1
135.15.13.3
135.15.35.3
135.15.34.3
135.15.23.3
4.4.4.4
135.15.104.1
135.15.10.4
135.15.14.4
135.15.20.4
135.15.34.4
5.5.5.5
135.15.105.1
135.15.20.5
135.15.25.5
135.15.35.5
6.6.6.6
135.15.106.1
135.15.26.6
135.15.16.6
7.7.7.7
135.15.107.1
135.15.10.7
135.15.20.7
10.10.10.10
135.15.110.1
135.15.20.10
20.20.20.20
135.15.120.1
135.15.20.20
} {ping $address}
tclquit
```

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

12. SNMP Security

Issue: On R1, create an SNMP view test via SNMP of all objects in the system group except for system 7 objects. Also, add all objects of the Cisco private MIB.

Solution:

R1

```
snmp-server view test system included
snmp-server view test cisco included
snmp-server view test system.7 excluded
```

Configuration lines are used to create an SNMP view test of all objects in the system group except for system 7 objects, as well as all objects of Cisco private MIBs.

Issue: Configure read-only access for the operator from VLAN 20 only.
Configure read-write access for the administrator from VLAN 30 only.

Solution:

R1

```
access-list 86 remark --------- administrator SNMP access ---------
access-list 86 permit 135.15.20.0 0.0.0.255
access-list 86 remark -----------------------------------------------
access-list 87 remark --------- operator SNMP access --------------
access-list 87 permit 135.15.26.0 0.0.0.255
access-list 87 remark -----------------------------------------------
```

Access list 87 defines the source of the allowed SNMP requests from the user **operator** (VLAN 20).

Access list 86 defines the source of the allowed SNMP requests from the user **administrator** (VLAN 30).

Issue: On R1, configure a user **administrator** in the SNMP group **administratorgrp** to be able to read and write to view test via SNMP, if the user passes SNMP authentication. Use Message Digest 5 (MD5) and the password test.

Configure the user **operator** in the SNMP group **operatorgrp** to be able to read the view test via SNMP. No SNMP authentication is required for the user operator. Configure SNMPv3 for this task.

Solution:

R1

```
snmp-server user operator operatorgrp v3 access 87
snmp-server user administrator administratorgrp v3 auth md5 test access 86
snmp-server group operatorgrp v3 noauth read test
snmp-server group administratorgrp v3 auth read test write test
```

Lines 1 and 3 are used to configure a user **operator** in the group **operatorgrp** to be able to read the view test. No authentication is required for the user operator. SNMP version 3 is used.

Lines 2 and 4 are used to configure a user **administrator** in the group **administratorgrp** to be able to read and write to view test, if the user passes authentication. Use MD5 and the password **test.** SNMPv3 is used.

The **snmp-server user administrator administratorgrp v3 auth md5 test access 86** command will not be listed in the show run. Verify the results with the following command:

```
R1#show snmp user

User name: operator
Engine ID: 800000090300AABBCC000100
storage-type: nonvolatile        active access-list: 87
Authentication Protocol: None
Privacy Protocol: None
```



```
Group-name: operatorgrp

User name: administrator
Engine ID: 800000090300AABBCC000100
storage-type: nonvolatile        active access-list: 86
Authentication Protocol: MD5
Privacy Protocol: None
Group-name: administratorgrp
```

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

13. IPv6 Addressing

Issue: Configure site-local IPv6 addresses according to the IPv6 diagram. Configure link-local IPv6 addresses only if necessary.

Solution:

IPv6 address configuration is done on an interface using the **ipv6 address** command followed by subnet and prefix length. On multipoint nonbroadcast interfaces, the mapping between IPv6 and corresponding Layer 2 addresses must be correct. The link-level address will be selected automatically based on the current router hardware. To prevent sudden changes to link-level addresses, specify the link-level address statically.

Ethernet interface example:

R3
```
interface Ethernet0/0
 ipv6 address FEC2::13:3/125
!
```

Serial interface example with link local address:

R3
```
interface Serial1/0
 ipv6 address FE80::3 link-local
 ipv6 address FEC2::34:3/125
!
```

VPN example with link local address:

R1
```
interface Tunnel12
 ipv6 address FE80::1 link-local
 ipv6 address FEC2::12:1/125
!
```
Full addressing configuration for this scenario:

R1
```
interface Tunnel12
 ipv6 address FE80::1 link-local
 ipv6 address FEC2::12:1/125
interface Tunnel14
 ipv6 address FE80::1 link-local
 ipv6 address FEC2::14:1/125
interface Ethernet0/1.10
```

```
 ipv6 address FEC2::13:1/125
interface Ethernet0/1.40
 ipv6 address FEC2::16:1/125
```

R2

```
interface Tunnel12
 ipv6 address FE80::2 link-local
 ipv6 address FEC2::12:2/125
interface Ethernet0/0.20
 ipv6 address FE80::2 link-local
 ipv6 address FEC2::26:2/125
interface Ethernet0/0.60
 ipv6 address FEC2::25:2/125
interface Serial1/1
 ipv6 address FE80::2 link-local
 ipv6 address FEC2::23:2/125
```

R3

```
interface Ethernet0/0
 ipv6 address FEC2::13:3/125
interface Serial1/0
 ipv6 address FE80::3 link-local
 ipv6 address FEC2::34:3/125
interface Serial1/1
 ipv6 address FE80::3 link-local
 ipv6 address FEC2::23:3/125
```

R4

```
interface Tunnel14
 ipv6 address FE80::4 link-local
 ipv6 address FEC2::14:4/125
interface Ethernet0/1
 ipv6 address FE80::4 link-local
 ipv6 address FEC2::38:4/125
interface Serial1/0
 ipv6 address FE80::4 link-local
 ipv6 address FEC2::34:4/125
```

R5

```
interface Ethernet0/0.60
 ipv6 address FEC2::25:5/125
interface Ethernet0/1
 ipv6 address FEC2::35:5/125
```

R6

```
interface Ethernet0/0.20
 ipv6 address FE80::6 link-local
 ipv6 address FEC2::26:6/125
interface Ethernet0/0.40
 ipv6 address FEC2::16:6/125
```

R8

```
interface Ethernet0/0
 ipv6 address FE80::5 link-local
 ipv6 address FEC2::38:5/125
```

Verify that the other end of the link is reachable:

```
R1#ping fec2::12:2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to FEC2::12:2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 17/17/18 ms
```

Verify that the other end of the link is reachable using link local addresses:


```
R1#ping FE80::2
Output Interface: Tunnel12
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to FE80::2, timeout is 2 seconds:
Packet sent with a source address of FE80::1%Tunnel12
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 17/17/17 ms
```

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

14. IPv6 Routing

Issue: Configure the OSPFv3 backbone area on VLAN 20.
R2 must be always the DR and have router ID (RID) 200.200.200.2.
OSPFv3 hello packets must be unicast.

Solution:

Restrictions limit the network type to nonbroadcast. To make R2 a DR, the R6 OSPF priority must be set to 0. A neighbor statement must be issued on the DR to connect to the spoke (DROTHER).

R2
```
interface Ethernet0/0.20
 description VLAN20
 encapsulation dot1Q 20
 ipv6 address FE80::2 link-local
 ipv6 address FEC2::26:2/125
 ipv6 ospf 100 area 0
 ipv6 ospf neighbor FE80::6
 ipv6 ospf network non-broadcast
!
ipv6 router ospf 100
 router-id 200.200.200.2
```

R6
```
interface Ethernet0/0.20
 description VLAN20
 encapsulation dot1Q 20
 ipv6 address FE80::6 link-local
 ipv6 address FEC2::26:6/125
 ipv6 ospf 100 area 0
 ipv6 ospf network non-broadcast
 ipv6 ospf priority 0
!
 ipv6 router ospf 100
```

Verify IPv6 OSPF neighbors:

```
R6#show ipv6 ospf neighbor

            OSPFv3 Router with ID (135.15.106.1) (Process ID 100)

Neighbor ID     Pri   State           Dead Time   Interface ID    Interface
200.200.200.2     1       FULL/DR                     00:00:19         17
Ethernet0/0.20
```

```
135.15.101.1              0       FULL/     -                   00:01:51        20
Ethernet0/0.40
```

Issue: Configure OSPFv3 Area 25 between R2 and R5 over VLAN 60, and on R5 on interface E0/1.
The OSPF process on R5 must not possess knowledge of routing information beyond its area.

Solution:

A totally stubby area will be the appropriate attribute for Area 25. This will ensure that R5 knows only about its own area and will have ::/0 toward the Area Border Router (ABR):

R2

```
interface Ethernet0/0.60
 description VLAN60
 encapsulation dot1Q 60
 ipv6 address FEC2::25:2/125
 ipv6 ospf 100 area 25
!
ipv6 router ospf 100
 router-id 200.200.200.2
 area 25 stub no-summary
!
```

R5

```
interface Ethernet0/0.60
 description VLAN60
 encapsulation dot1Q 60
 ipv6 address FEC2::25:5/125
 ipv6 ospf 100 area 25
!
interface Ethernet0/1
 ipv6 address FEC2::35:5/125
 ipv6 ospf 100 area 25
!
ipv6 router ospf 100
 area 25 stub
!
```

Verify the status of this area:

```
R2#show ipv6 ospf 100 | b Area 25
    Area 25
        Number of interfaces in this area is 1
        It is a stub area, no summary LSA in this area
        Generates stub default route with cost 1
        SPF algorithm executed 2 times
        Number of LSA 8. Checksum Sum 0x04AC1B
        Number of DCbitless LSA 0
        Number of indication LSA 0
        Number of DoNotAge LSA 0
        Flood list length 0
```

Verify that R5 does not see any routes except the default route:

```
R5#show ipv6 route ospf
IPv6 Routing Table - default - 6 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
       H - NHRP, I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
       IS - ISIS summary, D - EIGRP, EX - EIGRP external, NM - NEMO
       ND - ND Default, NDp - ND Prefix, DCE - Destination, NDr - Redirect
       O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
```



```
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, ls - LISP site
       ld - LISP dyn-EID, a - Application
OI  ::/0 [110/11]
     via FE80::A8BB:CCFF:FE00:200, Ethernet0/0.60
```

Issue: Configure OSPFv3 Area 16 on the VLAN 40 link between R1 and R6, and over the VPN between R1 and R2.
On VLAN 40, ensure that OSPF announces the other end of the link as the host route.

Solution:

VLAN 40 link restrictions require you to configure a point-to-multipoint network type. Neighbors will discover each other automatically. Each side will announce itself as a /128 route.

R1

```
interface Ethernet0/1.40
 encapsulation dot1Q 40
 ipv6 address FEC2::16:1/125
 ipv6 ospf 100 area 16
 ipv6 ospf network point-to-multipoint
!
interface Tunnel12
 ipv6 address FE80::1 link-local
 ipv6 address FEC2::12:1/125
 ipv6 ospf 100 area 16
!
ipv6 router ospf 100
!
```

R2

```
interface Tunnel12
 ipv6 address FE80::2 link-local
 ipv6 address FEC2::12:2/125
 ipv6 ospf 100 area 16
!
```

R6

```
interface Ethernet0/0.40
 encapsulation dot1Q 40
 ipv6 address FEC2::16:6/125
 ipv6 ospf 100 area 16
 ipv6 ospf network point-to-multipoint
!
```

Verify that endpoints are advertised as host routes:

```
R1#show ipv6 route ospf | s /128
O   FEC2::16:6/128 [110/10]
     via FE80::A8BB:CCFF:FE00:600, Ethernet0/1.40

R2#show ipv6 route ospf | s /128
O   FEC2::16:1/128 [110/1000]
     via FE80::1, Tunnel12
O   FEC2::16:6/128 [110/1010]
     via FE80::1, Tunnel12

R6#show ipv6 route ospf | s /128
O   FEC2::16:1/128 [110/10]
     via FE80::A8BB:CCFF:FE00:110, Ethernet0/0.40
```

Issue: Subnet 23:0/125 must be external to OSPF.
Control redistribution so that it is limited to the 23:0/125 subnet.

Solution:

Subnet 23:0/125 must be redistributed into OSPF.

R2

```
ipv6 access-list Subnet-23:0
 sequence 20 permit ipv6 FEC2::23:0/125 any
!
route-map Connected-To-OSPFv3 permit 10
 match ipv6 address Subnet-23:0
!
ipv6 router ospf 100
 redistribute connected metric 1 route-map Connected-To-OSPFv3
!
```

Verify that R2 generates an external link-state advertisement (LSA):

```
R2#show ipv6 ospf database external self-originate

            OSPFv3 Router with ID (200.200.200.2) (Process ID 100)

                Type-5 AS External Link States

  LS age: 79
  LS Type: AS External Link
  Link State ID: 0
  Advertising Router: 200.200.200.2
  LS Seq Number: 80000005
  Checksum: 0xC480
  Length: 44
  Prefix Address: FEC2::23:0
  Prefix Length: 125, Options: None
  Metric Type: 2 (Larger than any link state path)
  Metric: 1
```

Issue: Configure EIGRP between R1 and R4, R1 and R3, R3 and R4, and R4 and R8.

Solution:

The IPv6 EIGRP configuration is done on the interface level:

R1

```
interface Tunnel14
 ipv6 address FE80::1 link-local
 ipv6 address FEC2::14:1/125
 ipv6 eigrp 3
!
interface Ethernet0/1.10
 description VLAN10
 encapsulation dot1Q 10
ipv6 address FEC2::13:1/125
 ipv6 eigrp 3
!
ipv6 router eigrp 3
 eigrp router-id 135.15.101.1
!
```

R3

```
interface Ethernet0/0
 ipv6 address FEC2::13:3/125
```



```
 ipv6 eigrp 3
!
interface Serial1/0
 ipv6 address FE80::3 link-local
 ipv6 address FEC2::34:3/125
 ipv6 eigrp 3
!
ipv6 router eigrp 3
 eigrp router-id 135.15.103.1
!
```

R4

```
interface Tunnel14
 ipv6 address FE80::4 link-local
 ipv6 address FEC2::14:4/125
 ipv6 eigrp 3
!
interface Ethernet0/1
 ipv6 address FE80::4 link-local
 ipv6 address FEC2::38:4/125
 ipv6 eigrp 3
!
interface Serial1/0
 ipv6 address FE80::4 link-local
 ipv6 address FEC2::34:4/125
 ipv6 eigrp 3
!
ipv6 router eigrp 3
 eigrp router-id 135.15.104.1
!
```

R8

```
interface Ethernet0/0
 ipv6 address FE80::5 link-local
 ipv6 address FEC2::38:5/125
 ipv6 eigrp 3
!
ipv6 router eigrp 3
 eigrp router-id 135.15.110.1
!
```

Verify that EIGRP has successfully established neighbor relationships between routers. Here is an example on R4:

```
R4#show ipv6 eigrp neighbors
EIGRP-IPv6 Neighbors for AS(3)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q
Seq
                                                  (sec)         (ms)       Cnt
Num
2   Link-local address:     Tu14                     11 01:54:26   39  5000  0
14
    FE80::1
1   Link-local address:     Et0/1                    10 01:54:27    9   100  0
5
    FE80::5
0   Link-local address:     Se1/0                    12 01:55:04   12   100  0
11
    FE80::3
```

Verify that EIGRP routes are advertised. Here is an example on R1:

```
R1#show ipv6 route eigrp
IPv6 Routing Table - default - 16 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
```

```
       H - NHRP, I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
       IS - ISIS summary, D - EIGRP, EX - EIGRP external, NM - NEMO
       ND - ND Default, NDp - ND Prefix, DCE - Destination, NDr - Redirect
       O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, ls - LISP site
       ld - LISP dyn-EID, a - Application
D   FEC2::34:0/125 [90/2195456]
     via FE80::A8BB:CCFF:FE00:300, Ethernet0/1.10
D   FEC2::38:0/125 [90/2221056]
     via FE80::A8BB:CCFF:FE00:300, Ethernet0/1.10
```

Issue: Subnet 23:0/125 must be external to EIGRP.
Apply the configuration on R3.

Solution:

Subnet 23:0/125 must be redistributed into EIGRP.

R3

```
ipv6 router eigrp 3
 redistribute connected
!
```

Verify that R3 inserts the route into EIGRP topology:

```
R3#show ipv6 eigrp topology | s 23:0/125
P FEC2::23:0/125, 1 successors, FD is 2169856
        via Rconnected (2169856/0)
```

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

15. IPv6 Redistribution

Issue: Redistribute EIGRP into OSPF on R1 using a minimal nonzero metric.
Redistribute OSPF into EIGRP on R1.

Solution:

The minimal metric is 1. When redistributing IPv6, the **redistribute connected** command or the **include-connected** keyword also need to be used; otherwise in IPv6 redistribution, connected networks will not be redistributed between routing protocols.

R1

```
ipv6 router eigrp 3
 redistribute ospf 100 metric 1000 100 255 1 1500 include-connected
!
ipv6 router ospf 100
 redistribute eigrp 3 metric 1 include-connected
!
```


Issue: Restrict redistribution to exclude the 23:0/125 subnet.

Solution:

Apply filtering to redistribution:

R1

```
ipv6 prefix-list IPv6-redistr-filter seq 5 permit FEC2::23:0/125
!
route-map Connected-2-EIGRP162 permit 10
 match ip address Connected-2-EIGRP162
!
route-map IPv6-redistr-filter deny 10
 match ipv6 address prefix-list IPv6-redistr-filter
!
route-map IPv6-redistr-filter permit 20
!
ipv6 router eigrp 3
 redistribute ospf 100 metric 1000 100 255 1 1500 route-map IPv6-redistr-filter
include-connected
!
ipv6 router ospf 100
 redistribute eigrp 3 metric 1 route-map IPv6-redistr-filter include-connected
!
```

Issue: Ensure that all global IPv6 addresses specified in the IPv6 routing section are reachable.

Solution:

Verify that R3 inserts the route into the EIGRP topology and generates OSPF LSAs:

```
R1#show ipv6 eigrp topology
EIGRP-IPv6 Topology Table for AS(3)/ID(135.15.101.1)
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

P FEC2::16:6/128, 1 successors, FD is 2585600
        via Redistributed (2585600/0)
P FEC2::35:0/125, 1 successors, FD is 2585600
        via Redistributed (2585600/0)
P FEC2::26:0/125, 1 successors, FD is 2585600
        via Redistributed (2585600/0)
P FEC2::34:0/125, 1 successors, FD is 2195456
        via FE80::A8BB:CCFF:FE00:300 (2195456/2169856), Ethernet0/1.10
        via FE80::4 (27392000/2169856), Tunnel14
P FEC2::25:0/125, 1 successors, FD is 2585600
        via Redistributed (2585600/0)
P FEC2::23:0/125, 0 successors, FD is Infinity
        via FE80::A8BB:CCFF:FE00:300 (2195456/2169856), Ethernet0/1.10
        via FE80::4 (27904000/2681856), Tunnel14
P FEC2::12:0/125, 1 successors, FD is 2585600
        via Redistributed (2585600/0)
P FEC2::13:0/125, 1 successors, FD is 281600
        via Connected, Ethernet0/1.10
P FEC2::38:0/125, 1 successors, FD is 2221056
        via FE80::A8BB:CCFF:FE00:300 (2221056/2195456), Ethernet0/1.10
        via FE80::4 (26905600/281600), Tunnel14
P FEC2::14:0/125, 1 successors, FD is 26880000
        via Connected, Tunnel14
P FEC2::16:0/125, 1 successors, FD is 2585600
        via Redistributed (2585600/0)

R1#show ipv6 ospf database external self-originate  | i (^$)|(^   )|(Prefix)

            OSPFv3 Router with ID (135.15.101.1) (Process ID 100)
```

```
  Prefix Address: FEC2::13:0
  Prefix Length: 125, Options: None

  Prefix Address: FEC2::14:0
  Prefix Length: 125, Options: None

  Prefix Address: FEC2::34:0
  Prefix Length: 125, Options: None

  Prefix Address: FEC2::38:0
  Prefix Length: 125, Options: None
```

Below is a Tcl script to test universal reachability. To use the script, enter the **tclsh** command in privileged mode and paste this script. To kill failing pings, press **Ctrl-Shift** while pressing **6** twice. When you are done, enter the **tclquit** command to leave Tcl mode.

```
tclsh
foreach address {
FEC2::12:1
FEC2::13:1
FEC2::14:1
FEC2::16:1
FEC2::12:2
FEC2::23:2
FEC2::25:2
FEC2::26:2
FEC2::13:3
FEC2::23:3
FEC2::34:3
FEC2::14:4
FEC2::34:4
FEC2::38:4
FEC2::25:5
FEC2::35:5
FEC2::16:6
FEC2::26:6
FEC2::38:5
} {ping $address}
tclquit
```

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

16. QoS

Issue: Traffic between the networks 135.15.103.0/24 and 135.15.106.0/24 should be marked with IP precedence 2.
The marking should be retained throughout the network.
Accomplish this without the use of route maps.

Solution:

You have two options to accomplish this task:

- Configure committed access rate (CAR).
- Supply a Modular QoS CLI (MQC) configuration.

Both of these methods allow you to set the IP precedence bits without using a route map. For further guidance on how to fulfill this configuration requirement, read the Traffic Control section presented earlier.

Router R3 has MQC configured, and router R6 has rate limit configured to mark IP traffic between the mentioned networks with precedence 2.

R3

```
ip access-list extended QOS
 permit ip 135.15.103.0 0.0.0.255 135.15.106.0 0.0.0.255
!
class-map match-all QOS
 match access-group name QOS
!
policy-map QOS
 class QOS
  set precedence 2
!
interface Serial1/1
 service-policy output QOS
!
```

R6

```
access-list 101 remark --------- QOS ---------
access-list 101 permit ip 135.15.106.0 0.0.0.255 135.15.103.0 0.0.0.255
access-list 101 remark ----------------------
!
interface Ethernet0/0.20
 rate-limit output access-group 101 100000000 18750000 25000000 conform action
set-prec-transmit 2 exceed-action set-prec-transmit 2
!
```

Issue: Configure R2 and mark IP traffic coming from R3 with flash precedence, unless it is already marked with immediate precedence.

Solution:

This requires configuration of policy to match inbound traffic and remark it.

R2

```
class-map match-all FROM-R3
 match not ip precedence 2
!
policy-map FROM-R3
 class FROM-R3
  set precedence 3
!
interface Serial1/1
 service-policy input FROM-R3
!
```

Verify QoS operation:

```
R2#show policy-map interface Serial1/1
 Serial1/1

  Service-policy input: FROM-R3

    Class-map: FROM-R3 (match-all)
      910 packets, 202210 bytes
      5 minute offered rate 0000 bps, drop rate 0000 bps
      Match: not ip precedence 2
      QoS Set
```

```
        precedence 3
          Packets marked 910

    Class-map: class-default (match-any)
      0 packets, 0 bytes
      5 minute offered rate 0000 bps, drop rate 0000 bps
      Match: any
```

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

17. Switching Specialties

Issue: On SW4, configure ports, 0/3, 1/0, and 1/1 in access mode in VLAN 30. On each of these three interfaces, only one MAC address is allowed.

Solution:

Configure port security by first enabling port security, then by configuring the maximum number of MAC addresses allowed:

SW4:

```
int Ethernet0/3
 switchport access vlan 30
 switchport mode access
 switchport port-security
 switchport port-security max 1
!
int Ethernet1/0
 switchport access vlan 30
 switchport mode access
 switchport port-security
 switchport port-security max 1
!
int Ethernet1/1
 switchport access vlan 30
 switchport mode access
 switchport port-security
 switchport port-security max 1
!
```

Issue: What should the switch do when the number of port-secure MAC addresses reaches the maximum limit allowed on the port?

Solution:

The switch may perform one of three actions:

- Only drop packets with unknown source addresses on the Ethernet1/1 interface of SW4:

```
switchport port-security violation protect
```

- Drop packets with unknown source addresses and notify on the Ethernet1/0 interface of SW4:


```
switchport port-security violation restrict
```

- Shut down the Ethernet0/3 interface of SW4:

```
switchport port-security violation shutdown
```

The default behavior is to shut down the interface.

```
SW4#show port-security
Secure Port  MaxSecureAddr  CurrentAddr  SecurityViolation  Security Action
                (Count)       (Count)          (Count)
---------------------------------------------------------------------------
      Et0/3              1            0                  0         Shutdown
      Et1/0              1            0                  0          Protect
      Et1/1              1            0                  0         Restrict
---------------------------------------------------------------------------
Total Addresses in System (excluding one mac per port)     : 0
Max Addresses limit in System (excluding one mac per port) : 4096
```

Issue: Port 0/3 should be automatically brought out of the error-disabled state after a time-out of 50 seconds.

Solution:

Enable error recovery in global configuration mode:

```
errdisable recovery cause psecure-violation
errdisable recovery interval 50
```

Below is a quote from the Cisco Command Reference about the **errdisable recovery** command.

> "A cause (link-flap, bpduguard, and so forth) is defined as the reason that the error-disabled state occurred. When a cause is detected on a port, the port is placed in the error-disabled state, an operational state similar to the link-down state.
>
> When a port is error-disabled, it is effectively shut down, and no traffic is sent or received on the port. For the bridge protocol data unit (BPDU) guard and port-security features, configure the switch to shut down just the offending VLAN on the port when a violation occurs, instead of shutting down the entire port.
>
> If you do not enable the recovery for the cause, the port stays in the error-disabled state until you enter the **shutdown** and the **no shutdown** interface configuration commands. If you enable the recovery for a cause, the port is brought out of the error-disabled state and allowed to retry the operation again when all the causes have timed out.
>
> Otherwise, you must enter the **shutdown** and then the **no shutdown** commands to manually recover a port from the error-disabled state."

Verify **errdisable recovery** configuration:

```
SW4#show errdisable recovery | e Dis|
Recovery Status                 Timer Status
---------------                 ------------
udld                              Disabled
bpduguard                         Disabled
security-violation                Disabled
channel-misconfig                 Disabled
vmps                              Disabled
pagp-flap                         Disabled
dtp-flap                          Disabled
link-flap                         Disabled
```

```
l2ptguard                                Disabled
psecure-violation                        Enabled
gbic-invalid                             Disabled
dhcp-rate-limit                          Disabled
mac-limit                                Disabled
unicast-flood                            Disabled
storm-control                            Disabled
arp-inspection                           Disabled
loopback                                 Disabled
link-monitor-failure                     Disabled
oam-remote-failure critical-event        Disabled
oam-remote-failure dying-gasp            Disabled
oam-remote-failure link-fault            Disabled
dot1ad-incomp-etype                      Not supported
dot1ad-incomp-tunnel                     Not supported
mvrp                                     Not supported
transceiver-incomp                       Not supported
inline-power                             Not supported

Timer interval: 50 seconds

Interfaces that will be enabled at the next timeout:
```

Issue: Configure port-security aging.

Solution:

Configure port security aging on ports 1/0 and 1/1 of SW4:

```
interface Ethernet1/0
 switchport port-security aging time 2
 switchport port-security aging type inactivity
!
interface Ethernet1/1
 switchport port-security aging time 120
!
```

By default, the port security aging feature is disabled.
The default time is 0 minutes.
The default aging type is absolute.

Below is a quote from the Cisco Command Reference about switchport port security aging.

> "To enable secure address aging for a particular port, set the aging time to a value other than 0 for that port.
>
> To allow limited-time access to particular secure addresses, set the aging type as absolute. When the aging time lapses, the secure addresses are deleted.
>
> To allow continuous access to a limited number of secure addresses, set the aging type as inactivity. This removes the secure address when it becomes inactive, and other addresses can become secure."

Verify port security aging configuration:

```
SW4#show port-security interface e0/3
Port Security              : Enabled
Port Status                : Secure-up
Violation Mode             : Shutdown
Aging Time                 : 0 mins
Aging Type                 : Absolute
SecureStatic Address Aging : Disabled
```



```
Maximum MAC Addresses      : 1
Total MAC Addresses        : 0
Configured MAC Addresses   : 0
Sticky MAC Addresses       : 0
Last Source Address        : 0000.0000.0000
Last Source Address VlanId : 0
Security Violation Count   : 0

SW4#show port-security interface e1/0
Port Security              : Enabled
Port Status                : Secure-up
Violation Mode             : Protect
Aging Time                 : 2 mins
Aging Type                 : Inactivity
SecureStatic Address Aging : Disabled
Maximum MAC Addresses      : 1
Total MAC Addresses        : 0
Configured MAC Addresses   : 0
Sticky MAC Addresses       : 0
Last Source Address        : 0000.0000.0000
Last Source Address VlanId : 0
Security Violation Count   : 0

SW4#show port-security interface e1/1
Port Security              : Enabled
Port Status                : Secure-up
Violation Mode             : Restrict
Aging Time                 : 120 mins
Aging Type                 : Absolute
SecureStatic Address Aging : Disabled
Maximum MAC Addresses      : 1
Total MAC Addresses        : 0
Configured MAC Addresses   : 0
Sticky MAC Addresses       : 0
Last Source Address        : 0000.0000.0000
Last Source Address VlanId : 0
Security Violation Count   : 0
```

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

18. System Administration

Issue: The network administrator would like to have the Telnet session to R6 on port 3004 redirected to R2 on port 3007.

Solution:

When you want to use Telnet to go to R6, use the following command:

```
telnet <R6_IP _ADDRESS> 3004
```

Port 3000 is a base Telnet TCP port for rotaries. If you want to use Telnet to go to vty line number 6, use port 3006. The redirection to R2's vty line 4 on port 3007 can be done with the autocommand **telnet** *R2_ip_address* **3007** under the vty line 4 of the router R6.

R2

```
line vty 4
 rotary 7
!
```

R6

```
line vty 4
 rotary 4
 autocommand  telnet 135.15.102.1 3007
!
```

Issue: Redirection should occur only if user **admin** with the password **test** is authenticated on routers R6 and R2.

Solution:

On routers R2 and R6, configure **username admin password test** and perform the local authentication under vty line 4, with the **login local** command.

R2

```
       username admin password 0 test
!
line vty 4
 login local
!
```

R6

```
username admin password 0 test
!
line vty 4
 login local
!
```

Issue: Allow a Telnet session to R6's port 3004 only from the 135.35.101.0/24 network.

Solution:

Configure an access list permitting network 135.35.101.0/24 and apply **access-class in** under the VTY line 4 of router R6.

R6

```
access-list 6 remark -------- Router Access Management ---------
access-list 6 permit 135.15.101.0 0.0.0.255
access-list 6 remark ------------------------------------------
!
line vty 4
 access-class 6 in
!
```

Verification:

Try to use Telnet to go to the R6 vty line 4 on port 3004 and verify that you can reach R2 after twice typing the username **admin** and password **test**:

```
R1#telnet 135.15.16.6 3004 /source-interface loo101
Trying 135.15.16.6, 3004 ... Open

------------------------------------------------
Cisco 360 R&S Workbook Labs
Product, POD location:  cierswbv5-ce-lab03-sc, SJ
Device:                 R6
------------------------------------------------
```



```
User Access Verification

Username: admin
Password: Trying 135.15.102.1, 3007 ... Open

------------------------------------------------
Cisco 360 R&S Workbook Labs
Product, POD location:  cierswbv5-ce-lab03-sc, SJ
Device:                 R2
------------------------------------------------

User Access Verification

Username: admin
Password:
R2>exit

[Connection to 135.15.102.1 closed by foreign host]
[Connection to 135.15.16.6 closed by foreign host]
```

Issue: The network administrator would like to have ability to execute commands on R5 using remote shell protocol.
User TESTADMIN from R4 should be able to see R5's output of any commands, for example, the **show run** and **show version** commands.
User TESTOPERATOR from R4 should be able to see R5's output of nonprivileged commands, for example, the **show version** command but not the **show run** command.

Solution:

The desired solution is based on using remote shell (rsh) and remote command (rcmd) protocols.

You can configure the following commands on router R5:

```
ip rcmd rsh-enable
ip rcmd remote-host TESTADMIN 135.15.20.4 R4 enable
ip rcmd remote-host TESTOPERATOR 135.15.20.4 R4
```

Then, from R4 you can run the commands on R5:

```
R4#rsh 135.15.20.5 /user TESTOPERATOR sh ver
Cisco IOS Software, Linux Software (I86BI_LINUX-ADVENTERPRISEK9-M), Version
15.3(1.3)T, ENGINEERING WEEKLY BUILD, synced to V152_4_M1_10
Copyright (c) 1986-2012 by Cisco Systems, Inc.
Compiled Thu 25-Oct-12 04:35 by hlo
...   skipped ...
```

The **show run** command is not allowed for the user TESTOPERATOR.

```
R4#rsh 135.15.20.5 /user TESTOPERATOR sh run

Line has invalid autocommand "sh run"
```

The **show run** command is allowed for TESTADMIN because this user can run commands in enable mode:

```
R4#rsh 135.15.20.5 /user TESTADMIN sh run

Building configuration...

Current configuration : 3429 bytes
!
version 15.3
service timestamps debug datetime msec
```

```
service timestamps log datetime msec
no service password-encryption
!
hostname R5
!
...   skipped ...
```

Issue: Router R5 should send a notification to indicate that free processor memory has fallen below 80,000 KB.

Solution:

The Memory Threshold Notifications feature allows you to reserve memory for critical notifications and to configure a router to issue notifications when available memory falls below a specified threshold. This feature was introduced in Cisco IOS Software Release 12.3(4)T and is documented at:

Basic System Management Configuration Guide, Cisco IOS Release 15M&T
- Memory Threshold Notifications

Set the processor memory threshold to 80,000 KB.
When available processor memory falls below this threshold, a notification message is triggered.

```
memory free low-watermark processor 80000
```

Issue: Reserve 2500 KB of memory so that critical operations such as event logging will continue to function even when router memory is exhausted.

Solution:

Reserve memory for use by critical processes:

```
memory reserve critical 2500
```

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

19. Multicast

Issue: Configure Protocol Independent Multicast (PIM) sparse mode (SM) routing on all routers involved in this scenario.

Multicast Diagram

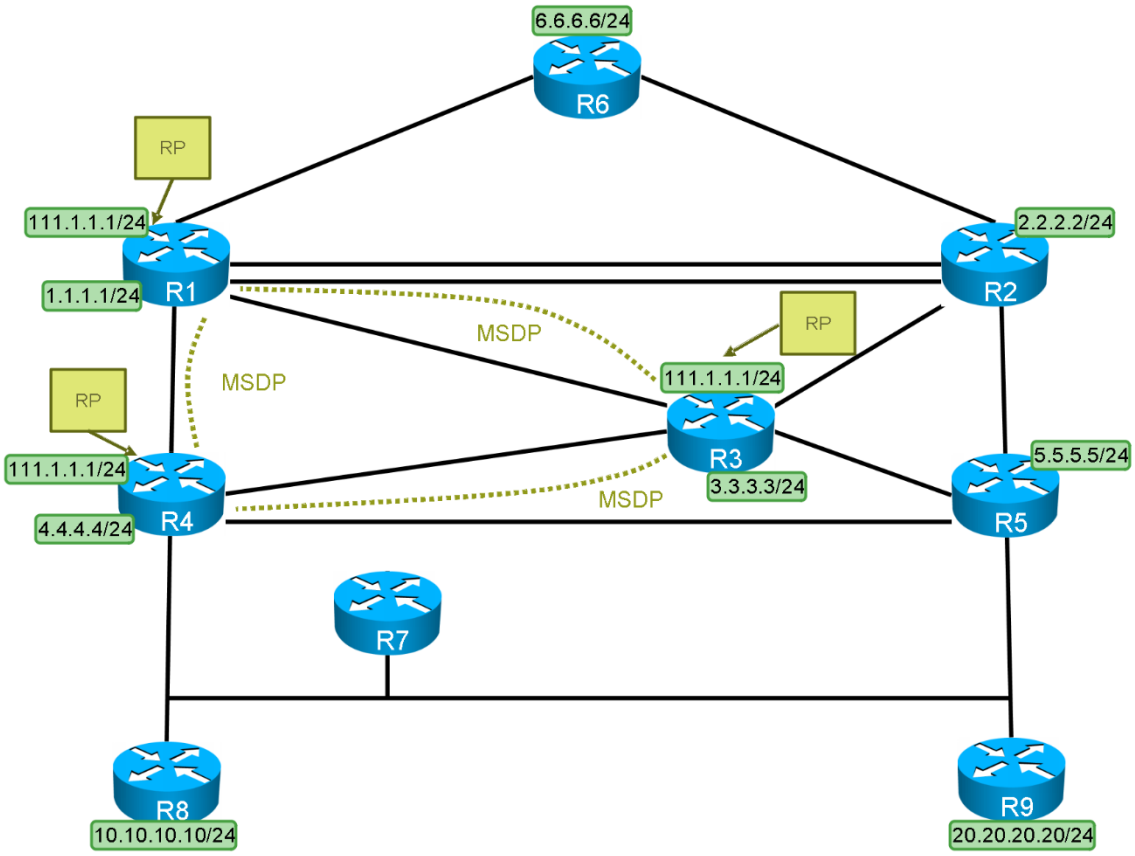

Solution:

Use the **ip pim sparse-mode** interface configuration command on all interfaces required in this scenario.

R1

```
interface Loopback1
 ip pim sparse-mode
interface Loopback10
 ip pim sparse-mode
interface Tunnel12
 ip pim sparse-mode
interface Tunnel14
 ip pim sparse-mode
interface Ethernet0/1.10
 ip pim sparse-mode
interface Ethernet0/1.40
 ip pim sparse-mode
```

R2

```
interface Loopback2
 ip pim sparse-mode
interface Tunnel12
 ip pim sparse-mode
interface Ethernet0/0.20
 ip pim sparse-mode
interface Ethernet0/0.60
 ip pim sparse-mode
interface Serial1/1
 ip pim sparse-mode
```

R3

```
interface Loopback3
 ip pim sparse-mode
interface Loopback10
 ip pim sparse-mode
interface Ethernet0/0
 ip pim sparse-mode
interface Ethernet0/1
 ip pim sparse-mode
interface Serial1/0
 ip pim sparse-mode
interface Serial1/1
 ip pim sparse-mode
```

R4

```
interface Loopback4
 ip pim sparse-mode
interface Loopback10
 ip pim sparse-mode
interface Tunnel10
 ip pim sparse-mode
interface Tunnel14
 ip pim sparse-mode
interface Ethernet0/0
 ip pim sparse-mode
interface Serial1/0
 ip pim sparse-mode
```

R5

```
interface Loopback5
 ip pim sparse-mode
interface Ethernet0/0.30
 ip pim sparse-mode
interface Ethernet0/0.60
 ip pim sparse-mode
interface Ethernet0/1
 ip pim sparse-mode
```

R6

```
interface Loopback6
 ip pim sparse-mode
interface Ethernet0/0.20
 ip pim sparse-mode
interface Ethernet0/0.40
 ip pim sparse-mode
```

R7

```
interface Loopback7
 ip pim sparse-mode
interface Tunnel10
 ip pim sparse-mode
interface Ethernet0/0
 ip pim sparse-mode
```


R8

```
interface Loopback10
 ip pim sparse-mode
interface Ethernet0/0
 ip pim sparse-mode
```

R9

```
interface Loopback20
 ip pim sparse-mode
interface Ethernet0/0
ip pim sparse-mode
```

Issue: Join multicast group 225.22.22.22 on the loopback interfaces.

Solution:

IP connectivity to loopback interfaces is provided by the BGP routing process. Use the **ip igmp join-group 225.22.22.22** interface configuration command on each loopback interface involved in this scenario. Also, enable PIM on these loopbacks.

R1

```
interface Loopback1
 ip address 1.1.1.1 255.255.255.0
 ip pim sparse-mode
 ip igmp join-group 225.22.22.22
```

R2

```
interface Loopback2
 ip address 2.2.2.2 255.255.255.0
 ip pim sparse-mode
 ip igmp join-group 225.22.22.22
```

R3

```
interface Loopback3
 ip address 3.3.3.3 255.255.255.0
 ip pim sparse-mode
 ip igmp join-group 225.22.22.22
```

R4

```
interface Loopback4
 ip address 4.4.4.4 255.255.255.0
 ip pim sparse-mode
 ip igmp join-group 225.22.22.22
```

R5

```
interface Loopback5
 ip address 5.5.5.5 255.255.255.0
 ip pim sparse-mode
 ip igmp join-group 225.22.22.22
```

R6

```
interface Loopback6
 ip address 6.6.6.6 255.255.255.0
 ip pim sparse-mode
 ip igmp join-group 225.22.22.22
```

R7

```
interface Loopback7
 ip address 7.7.7.7 255.255.255.0
 ip pim sparse-mode
 ip igmp join-group 225.22.22.22
```

R8

```
interface Loopback10
 ip address 10.10.10.10 255.255.255.0
 ip pim sparse-mode
 ip igmp join-group 225.22.22.22
```

R9

```
interface Loopback20
 ip address 20.20.20.20 255.255.255.0
 ip pim sparse-mode
 ip igmp join-group 225.22.22.22
```

Issue: Configure R1, R3, and R4 as redundant RPs in the network. Use 111.1.1.1/32 as the RP address. Use a static configuration on all routers and switches that need to access this RP address. Configure all three routers to cache source/group pairs to reduce join latency.

Solution:

Configure anycast RP. Routers R1, R3, and R4 will be configured with the same loopback IP address, 111.1.1.1/32.

R1

```
interface Loopback10
 ip address 111.1.1.1 255.255.255.255
 ip pim sparse-mode
```

R3

```
interface Loopback10
 ip address 111.1.1.1 255.255.255.255
 ip pim sparse-mode
```

R4

```
interface Loopback10
 ip address 111.1.1.1 255.255.255.255
 ip pim sparse-mode
```

All routers in the multicast network will be configured with the **ip pim rp-address 111.1.1.1** command. Multicast routers will use the closest RP by performing an IGP lookup for the 111.1.1.1/32 address.

R1

```
ip pim rp-address 111.1.1.1
```

R2

```
ip pim rp-address 111.1.1.1
```

R3

```
ip pim rp-address 111.1.1.1
```

R4

```
ip pim rp-address 111.1.1.1
```

R5

```
ip pim rp-address 111.1.1.1
```

R6

```
ip pim rp-address 111.1.1.1
```

R7

```
ip pim rp-address 111.1.1.1
```

R8


```
ip pim rp-address 111.1.1.1
```

R9

```
ip pim rp-address 111.1.1.1
```

The Multicast Source Discovery Protocol (MSDP) is fundamental for the anycast RP configuration. Configure MSDP peer relationships and mesh groups between R1, R3, and R4.

MSDP peers exchange Source-Active (SA) messages. Once the router forwards the MSDP SA information, it does not store it in memory. Therefore, if a member joins a group soon after an SA message is received by the local RP, that member will need to wait until the next SA message to hear about the source. This delay is known as join latency.

Using the **ip msdp cache-sa-state** command, you can configure a router to cache source/group pairs to reduce join latency in exchange to higher memory utilization.

To configure MSDP, use the following commands on all anycast RP routers. The following is an example configuration:

R1

```
ip msdp peer 135.15.103.1 connect-source Loopback101
ip msdp peer 135.15.104.1 connect-source Loopback101
ip msdp cache-sa-state
ip msdp originator-id Loopback101
ip msdp mesh-group RP 135.15.103.1
ip msdp mesh-group RP 135.15.104.1
```

R3

```
ip msdp peer 135.15.101.1 connect-source Loopback103
ip msdp peer 135.15.104.1 connect-source Loopback103
ip msdp cache-sa-state
ip msdp originator-id Loopback103
ip msdp mesh-group RP 135.15.101.1
ip msdp mesh-group RP 135.15.104.1
```

R4

```
ip msdp peer 135.15.101.1 connect-source Loopback104
ip msdp peer 135.15.103.1 connect-source Loopback104
ip msdp cache-sa-state
ip msdp originator-id Loopback104
ip msdp mesh-group RP 135.15.101.1
ip msdp mesh-group RP 135.15.103.1
```

Verify MSDP peering. The following is an example from router R4**:**

```
R4#show ip msdp summary
MSDP Peer Status Summary
Peer Address     AS    State    Uptime/  Reset SA    Peer Name
                                Downtime Count Count
135.15.101.1     126   Up       11:58:18 0     0     ?
135.15.103.1     ?     Up       11:58:18 0     0     ?
```

Verify MSDP SA state caching. The following is an example from router R4 with R7 actively sending a multicast stream:

```
R4#show ip msdp sa-cache
MSDP Source-Active Cache - 1 entries
(135.15.20.7, 225.22.22.22), RP 135.15.103.1, BGP/AS 0, 00:00:23/00:05:41, Peer
135.15.103.1
```

Issue: Each router and switch should receive a reply from every router when the address 225.22.22.22 is pinged.

Solution:

Perform a ping of multicast IP address from each router and verify that the replies are received. This is verification example from R1:

```
R1#ping 225.22.22.22
Type escape sequence to abort.
Sending 1, 100-byte ICMP Echos to 225.22.22.22, timeout is 2 seconds:

Reply to request 0 from 1.1.1.1, 13 ms
Reply to request 0 from 20.20.20.20, 61 ms
Reply to request 0 from 7.7.7.7, 53 ms
Reply to request 0 from 20.20.20.20, 53 ms
Reply to request 0 from 2.2.2.2, 51 ms
Reply to request 0 from 2.2.2.2, 51 ms
Reply to request 0 from 10.10.10.10, 48 ms
Reply to request 0 from 10.10.10.10, 48 ms
Reply to request 0 from 20.20.20.20, 44 ms
Reply to request 0 from 2.2.2.2, 39 ms
Reply to request 0 from 10.10.10.10, 39 ms
Reply to request 0 from 7.7.7.7, 39 ms
Reply to request 0 from 10.10.10.10, 35 ms
Reply to request 0 from 5.5.5.5, 35 ms
Reply to request 0 from 4.4.4.4, 34 ms
Reply to request 0 from 7.7.7.7, 31 ms
Reply to request 0 from 20.20.20.20, 29 ms
Reply to request 0 from 1.1.1.1, 29 ms
Reply to request 0 from 2.2.2.2, 28 ms
Reply to request 0 from 1.1.1.1, 19 ms
Reply to request 0 from 6.6.6.6, 19 ms
Reply to request 0 from 3.3.3.3, 19 ms
Reply to request 0 from 1.1.1.1, 19 ms
Reply to request 0 from 6.6.6.6, 17 ms
Reply to request 0 from 6.6.6.6, 17 ms
Reply to request 0 from 6.6.6.6, 17 ms
Reply to request 0 from 1.1.1.1, 13 ms
```

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

