360RSW05-WKB-CA07

Cisco 360 CCIE R&S Exercise Workbook Introduction

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

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

Cisco 360 CCIE R&S Exercise Workbook Lab 7 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 7 Configuration Section Answer Key

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

Grading and Duration

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

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

Difficulty Level

- Difficulty: Intermediate to Advanced

Restrictions and Goals

Note	Read this section carefully.

- To receive credit for a subsection, you must fully complete the subsection per the requirements. You will *not* receive partial credit for partially completed subsections.
- IPv4 subnets that are displayed in the scenario diagram belong to network 172.16.0.0/16.
- *Points will be deducted from multiple sections for failing to assign correct IPv4 addresses.*
- Do not use any static routes.
- Advertise loopback interfaces with their original masks for IPv4 protocol.
- Do not use the **ip default-network** or **ip default-gateway** commands.
- All IP addresses that are involved in the same virtual routing and forwarding (VRF) instance must be reachable, unless an explicitly stated filtering requirement restricts reachability.
- Do not introduce any new IPv4 addresses.
- Networks do not have to be reachable outside of their VRF.

- Use conventional routing algorithms only.
- Do not modify the hostname, console, or vty configuration unless you are specifically asked to do so.
- Do not modify the initial interface or IP address numbering.

Explanation of Each of the Restrictions and Goals

IPv4 subnets that are displayed in the scenario IPv4 IGP diagram belong to network 172.16.0.0/16.

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

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.

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 involved in the same VRF must be reachable, unless an explicitly stated filtering requirement restricts reachability.

This is a key goal to observe. This requires that all of your IGPs are configured properly. In addition, all of your routing policy tasks must 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 throughout this lab is that the term "redistribution" is never explicitly used. However, you must perform redistribution to assure that all IP addresses are reachable without the use of static routes.

Networks do not have to be reachable outside of their VRF.

This statement relaxes the requirement above. IP addresses in the CustomerA VRF do not have to be reachable from IP addresses in the default, or native, VRF.

Use conventional routing algorithms.

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

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

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

1. Switch Configuration

General Tasks:

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

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

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

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

See the following diagram for the VLAN layout.

VLAN Propagation Diagram

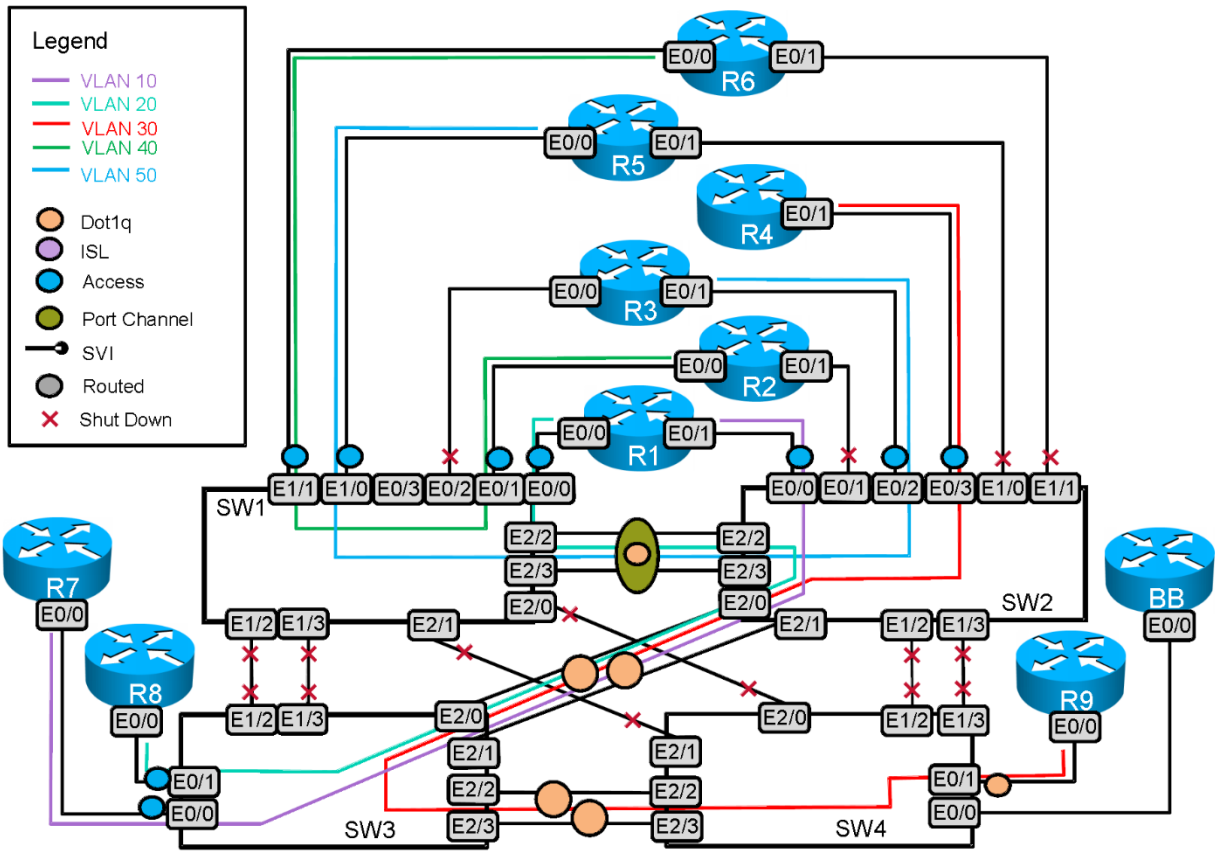

Issue: Do not advertise VLAN database information between switches.

Solution:

Set the required VTP mode: The task hints that you should look around at the entire exam to determine which VTP mode to use. VLAN database information must not be exchanged between switches. So you need to set the VTP mode to transparent.

Issue: Configure VLANs, and switch port access VLAN interfaces.

The task requires specific VLAN names. Here is an example of VLAN 40 configuration on SW1:

```
SW1#show run vlan 40
Building configuration...

Current configuration:
!
vlan 40
name VLAN40
end
```

All other VLANs are configured similarly according to the VLAN name tables.

Verification:

```
SW1#sh vlan brie

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active    Et0/2, Et0/3, Et1/2, Et1/3
                                                Et2/0, Et2/1
10   VLAN10                           active
20   VLAN20                           active    Et0/0
40   VLAN40                           active    Et0/1, Et1/1
50   VLAN50                           active    Et1/0
1002 fddi-default                     act/unsup
1003 token-ring-default               act/unsup
1004 fddinet-default                  act/unsup
1005 trnet-default                    act/unsup
SW1#
```

When assigning ports to access VLANs on Cisco Catalyst switches, specify the access mode. If you leave the trunk mode to the default *dynamic desirable*, the ports sometimes fail to link properly. It is also a good idea to add descriptions while the information is fresh in your mind.

Here is an example on SW2:

```
interface Ethernet0/0
 switchport access vlan 10
 switchport mode access
```

After assigning your access ports to the required VLANs, issue the commands **show VLAN** and **show interface status** to check your work. The latter command gives a one-line summary of each interface. Here is an example on SW1:

```
SW1#show interface status

Port         Name               Status       Vlan       Duplex  Speed Type
Et0/0                           connected    20           auto   auto unknown
Et0/1                           connected    40           auto   auto unknown
Et0/2                           disabled     1            auto   auto unknown
Et0/3        Visitor Connectivi connected    1            auto   auto unknown
Et1/0                           connected    50           auto   auto unknown
```

```
Et1/1                          connected    40          auto   auto unknown
Et1/2                          disabled     1           auto   auto unknown
Et1/3                          disabled     1           auto   auto unknown
Et2/0                          disabled     1           auto   auto unknown
Et2/1                          disabled     1           auto   auto unknown
Et2/2                          connected    trunk       auto   auto unknown
Et2/3                          connected    trunk       auto   auto unknown
Po1                            connected    trunk       auto   auto
SW1#
```

Verify that each router interface in each VLAN can ping the other router interfaces in that VLAN. Have you used the **no shut** command where required? Many candidates lose energy and time troubleshooting routing protocols when the real problem is basic connectivity within the subnet.

Issue: Configure a link between SW1 and SW2 with the aggregate bandwidth 20 Mb/s. Use ports Et2/2 and Et2/3 to accomplish this task.

Solution:

This can be accomplished by configuring EtherChannel between SW1 and SW2 using the parallel links 2/2 and 2/3. You can choose to negotiate the channel using PAgP (**desirable** and **auto**) or LACP (**active** and **passive**), or you can statically configure them with the mode set to **on**. The PAgP negotiation configuration is chosen in this answer key.

Interfaces in the channel must have identical configuration for ports in the channel. Remote ports must match the configuration of the local ports in the channel.

When the channel is configured, it's a good practice to configure channel port under interface Po1 (channel group is 1 in this example) for trunking. Commands from the Po1 interface will be replicated to all ports in the channel group automatically, and this way the configuration will match between ports and will be less error-prone.

SW1:

```
interface Ethernet2/2
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 20,50
 switchport mode trunk
 duplex auto
 channel-group 1 mode desirable
!
interface Ethernet2/3
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 20,50
 switchport mode trunk
 duplex auto
 channel-group 1 mode desirable
!
```

SW2:

```
interface Ethernet2/2
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 20,50
 switchport mode trunk
 duplex auto
 channel-group 1 mode auto
!
interface Ethernet2/3
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 20,50
 switchport mode trunk
 duplex auto
 channel-group 1 mode auto
!
```


Verification:

Here you see the output from the command **show etherchannel summary** indicating a working EtherChannel.

```
SW1#show etherchannel summary
Flags:  D - down        P - in port-channel
        I - stand-alone s - suspended
        H - Hot-standby (LACP only)
        R - Layer3      S - Layer2
        u - unsuitable for bundling
        U - in use      f - failed to allocate aggregator
        d - default port

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: Configure a dot1q trunk on all trunks involved in this scenario. Allow the minimal subset of VLANs on these trunks to satisfy the scenario.

Configuration:

The following sequence sets the encapsulation type to dot1q and the trunk mode to on. This example is on the Po1 interface:

```
interface Po1
  switchport trunk encapsulation dot1q
 switchport mode trunk
end
```

By default, traffic from all VLANs on a switch is allowed across trunk ports. Restricting this traffic is called “manually pruning” the VLANs on the trunk. It is generally considered a good practice to limit VLANs on a trunk to only those required, so that you limit unnecessary traffic and limit the extent of the spanning tree. The **switchport trunk allowed vlan** command is used to accomplish this.

```
SW2(config)#interface Po1
SW2(config-if)#switch trunk allowed vlan ?
  WORD    VLAN IDs of the allowed VLANs when this port is in trunking mode
  add     add VLANs to the current list
  all     all VLANs
  except  all VLANs except the following
  none    no VLANs
  remove  remove VLANs from the current list
```

Notice in the syntax above that the **add**, **except**, and **remove** keywords can be used in a very flexible way to define the allowed VLANs. If you just list the VLANs after the **allowed** keyword, it is an exclusive list. Note that the list of VLANs consists of comma-separated values, without spaces, and can include ranges. Note also that you can list VLANs that have not yet been created on the switch.

Verification:

There are many commands that will allow you to see the status of your trunks. We'll use **show interfaces trunk**:

```
SW2#show interfaces trunk

Port                Mode         Encapsulation  Status        Native vlan
Et2/0               on           802.1q         trunking      1
Et2/1               on           802.1q         trunking      1
Po1                 on           802.1q         trunking      1

Port                Vlans allowed on trunk
Et2/0               10,20,30
Et2/1               10,20,30
Po1                 20,50

Port                Vlans allowed and active in management domain
Et2/0               10,20,30
Et2/1               10,20,30
Po1                 20,50

Port                Vlans in spanning tree forwarding state and not pruned
Et2/0               10,20,30
Et2/1               10,20,30
Po1                 20,50
SW2#
```

The output above shows that the trunks are using dot1q encapsulation. It also indicates the VLANs that are allowed on the trunks, the VLANs that are allowed and actually created, and VTP pruning. By default, none are VTP pruned.

Remember to configure both sides of the trunk. Allowed VLANs should match. Try not to rely on trunk negotiation unless required. The native VLAN on the dot1q trunk must match at both ends.

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

2. IPv4 OSPF

Note	This section was partially preconfigured in the initialization file so that you could concentrate on more advanced topics.

Issue: Configure the OSPF backbone area between routers R1, R2, and R3.

Solution:

Area 0 is meant as an OSPF backbone area. OSPF Area 0 will have three routers connected to the subnet 172.16.123.0/24. We have three OSPF network types that do not elect a DR and BDR: point-to-point, point-to-multipoint, and point-to-multipoint non-broadcast. Since there are three devices on the subnet, point-to-point is not an option. There is no specific requirement for OSPF packet addressing (unicast or multicast), so either remaining OSPF network can be used. The **ip ospf network type point-to-multipoint** command is used in this answer key.

Verification:

Note that the point-to-multipoint OSPF network type models network as a collection of point-to-point interfaces, creating host routes for each address in the subnet. If you examine the routing table on R2, for example, you will see that there are /32 prefixes that represent the R1 and R3 interfaces on the 172.16.123.0/24 connected subnet.

```
R2#show ip route | inc 123\.[0-9]+\/
C        172.16.123.0/24 is directly connected, Ethernet0/2
O        172.16.123.1/32 [110/10] via 172.16.123.1, 03:43:07, Ethernet0/2
L        172.16.123.2/32 is directly connected, Ethernet0/2
O        172.16.123.3/32 [110/10] via 172.16.123.3, 03:43:02, Ethernet0/2
R2#

R2#show ip ospf interface brief
Interface    PID   Area            IP Address/Mask    Cost  State Nbrs F/C
Et0/2        1     0               172.16.123.2/24    10    P2MP  2/2
VL0          1     0               172.16.26.2/24     10    P2P   1/1
Lo251        1     2               172.16.25.1/30     1     LOOP  0/0
Lo255        1     2               172.16.25.5/30     1     LOOP  0/0
Lo259        1     2               172.16.25.9/30     1     LOOP  0/0
Et0/0        1     26              172.16.26.2/24     10    BDR   1/1
Lo102        1     101             172.16.102.1/24    1     P2P   0/0
R2#
```

Issue: Summarize /30 loopbacks in OSPF Area 2 to /27 into OSPF.

Solution:

Use the OSPF **area range** command to summarize OSPF intra-area prefixes. Note that the single best, longest match summary of these networks would be /28. Make sure that you are closely following the instructions.

```
R2#show running-config | sec router ospf
router ospf 1
 area 0 authentication message-digest
 area 2 range 172.16.25.0 255.255.255.224
 area 26 virtual-link 172.31.1.1 message-digest-key 1 md5 test
 summary-address 172.20.0.0 255.255.0.0
 redistribute connected subnets route-map Connected-->OSPF
 network 172.16.25.0 0.0.0.31 area 2
 network 172.16.26.0 0.0.0.255 area 26
 network 172.16.102.0 0.0.0.255 area 101
 network 172.16.123.0 0.0.0.255 area 0
R2#
```

Issue: Assign the IP address 172.20.10.1/24 to a loopback interface on R2 and add it into the OSPF routing process as an external major network.

Solution:

The meaning of external major networks is that the networks should be generated as an external LSA Type 5 network of E2 or E1 type. It also needs to be classful 172.20.0.0/16. The solution demonstrated in the Mentor Guide uses the command **redistribute connected subnets route-map Connected→OSPF**. Since all of the other connected networks are internal to OSPF, this route map is not strictly necessary. Then, a **summary-address** command is issued so that it is advertised with its classful mask.

R2:

```
router ospf 1
 area 0 authentication message-digest
 area 2 range 172.16.25.0 255.255.255.224
 area 26 virtual-link 172.31.1.1 message-digest-key 1 md5 test
 summary-address 172.20.0.0 255.255.0.0
 redistribute connected subnets route-map Connected-->OSPF
 network 172.16.25.0 0.0.0.31 area 2
 network 172.16.26.0 0.0.0.255 area 26
 network 172.16.102.0 0.0.0.255 area 101
 network 172.16.123.0 0.0.0.255 area 0
!
!
ip access-list standard Connected-->OSPF
 permit 172.20.10.0 0.0.0.255
!
!
route-map Connected-->OSPF permit 10
 match ip address Connected-->OSPF
!
!
```

Issue: Create a loopback on R6 and place it into OSPF Area 5

Solution:

If you take the time to draw a quick diagram, it is immediately apparent that OSPF Area 5 and Area 101 are not directly attached to OSPF Area 0. A virtual link is required across the transit OSPF Area 26. Remember that virtual links are configured between OSPF router IDs (RIDs), not interface addresses.

Issue: Authenticate the OSPF backbone area. Do not use a cleartext password.

Solution:

MD5 authentication will be used, because cleartext is specifically prohibited. When you authenticate Area 0, remember that the virtual link is an extension of Area 0, so Area 0 authentication should be configured on R6 as well.

Here is an example OSPF configuration on R6:

```
R6#show run | sec router ospf
router ospf 1
 area 0 authentication message-digest
 area 26 virtual-link 172.20.10.1 message-digest-key 1 md5 test
 network 172.16.26.0 0.0.0.255 area 26
 network 172.16.106.0 0.0.0.255 area 101
 network 172.31.1.1 0.0.0.0 area 5
R6#
```

The following output shows that authentication is enabled on the Ethernet0/2 interface and the OSPF virtual link:

```
R2#show ip ospf interface
```



```
Ethernet0/2 is up, line protocol is up
  Internet Address 172.16.123.2/24, Area 0, Attached via Network Statement
  Process ID 1, Router ID 172.20.10.1, Network Type POINT_TO_MULTIPOINT, Cost: 10
  Topology-MTID    Cost    Disabled    Shutdown      Topology Name
        0           10        no          no            Base
  Transmit Delay is 1 sec, State POINT_TO_MULTIPOINT
  Timer intervals configured, Hello 30, Dead 120, Wait 120, Retransmit 5
    oob-resync timeout 120
    Hello due in 00:00:27
  Supports Link-local Signaling (LLS)
  Cisco NSF helper support enabled
  IETF NSF helper support enabled
  Index 2/7, flood queue length 0
  Next 0x0(0)/0x0(0)
  Last flood scan length is 4, maximum is 5
  Last flood scan time is 0 msec, maximum is 1 msec
  Neighbor Count is 2, Adjacent neighbor count is 2
    Adjacent with neighbor 172.16.101.1
    Adjacent with neighbor 172.16.170.1
  Suppress hello for 0 neighbor(s)
  Message digest authentication enabled
    Youngest key id is 1
OSPF_VL0 is up, line protocol is up
  Internet Address 172.16.26.2/24, Area 0, Attached via Not Attached
  Process ID 1, Router ID 172.20.10.1, Network Type VIRTUAL_LINK, Cost: 10
  Topology-MTID    Cost    Disabled    Shutdown      Topology Name
        0           10        no          no            Base
  Configured as demand circuit
  Run as demand circuit
  DoNotAge LSA allowed
  Transmit Delay is 1 sec, State POINT_TO_POINT
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    oob-resync timeout 40
    Hello due in 00:00:09
  Supports Link-local Signaling (LLS)
  Cisco NSF helper support enabled
  IETF NSF helper support enabled
  Index 1/1, flood queue length 0
  Next 0x0(0)/0x0(0)
  Last flood scan length is 1, maximum is 1
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 172.31.1.1  (Hello suppressed)
  Suppress hello for 1 neighbor(s)
  Message digest authentication enabled
    Youngest key id is 1
```

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

3. IPv4 RIP

Issues: Configure RIPv2 over the connection between R1 and R4. Send RIP updates for a minimal number of routes from R1. Do not use summarization.

Solution:

A solution to the minimal number of routes problem for RIP is to originate a 0.0.0.0/0 prefix on router R1 by configuring the following command under the R1 RIP routing process: **default-**

information originate. With this command configured, R1 will advertise a 0.0.0.0/0 route to the downstream RIP routers. Unlike a summary, this technique does not suppress longer matches. To do so, we have filtered outbound updates using a prefix list. Only two networks are allowed in the prefix list: 0.0.0.0/0 and 172.16.101.0/24. The 172.16.101.0/24 subnet is used for the BGP peer relationship configuration on R4. Read the "MPLS Layer 3 VPNs" section.

R1:

```
router rip
 version 2
 passive-interface default
 no passive-interface Ethernet0/3
 network 172.16.0.0
 default-information originate
 distribute-list prefix RIP out Ethernet0/3
!
!
ip prefix-list RIP seq 5 permit 0.0.0.0/0
ip prefix-list RIP seq 10 permit 172.16.101.0/24
!
```

R4:

```
R4#show ip route rip
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       a - application route
       + - replicated route, % - next hop override

Gateway of last resort is 172.16.14.1 to network 0.0.0.0

R*    0.0.0.0/0 [120/1] via 172.16.14.1, 00:00:14, Ethernet0/3
      172.16.0.0/16 is variably subnetted, 11 subnets, 2 masks
R        172.16.101.0/24 [120/1] via 172.16.14.1, 00:00:14, Ethernet0/3
R4#
```

Issue: Make sure that RIP advertises only over the necessary interfaces.

Solution:

RIP uses a classful network statement to advertise the subnets of that major network. Therefore, the updates will be sent out from all interfaces matching the network statement. Use a passive interface to stop RIP from advertising out that interface. The recommended technique for placing the proper interfaces in a passive state is based upon a two-step configuration approach. First, enter the **passive-interface default** command under the RIP routing process. This step places all interfaces in a passive state. Then, selectively remove those interfaces that are required to run the protocol with the **no passive-interface** command.

Issue: Summarize routes when they are redistributed into OSPF.

Solution:

Use the **summary-address** command under the OSPF process to summarize the OSPF external prefixes that are redistributed from RIP. R1 will become an ASBR for the RIP redistributed prefixes.

R1:

```
router ospf 1
 area 0 authentication message-digest
 summary-address 172.16.192.0 255.255.252.0
 redistribute rip subnets route-map RIP-->OSPF
 passive-interface Tunnel10
 network 172.16.101.0 0.0.0.255 area 101
 network 172.16.123.0 0.0.0.255 area 0
!
!
route-map RIP-->OSPF permit 10
 match ip address RIP-->OSPF
!
!
!
ip access-list standard RIP-->OSPF
 permit 172.16.14.0
 permit 172.16.1.0
 permit 172.16.2.0
 permit 172.16.104.0
 permit 150.100.10.0
 permit 172.16.192.0 0.0.3.255
 permit 192.168.104.0 0.0.3.255
!
```

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

4. IPv4 EIGRP

Issue: Make sure that EIGRP AS 35 advertises only over the specified interfaces.

Solution:

This can be done in one of two ways: by using a passive interface or by using a network statement with a wildcard mask under the EIGRP process.

Verification:

To make sure you have met the requirements, check the output of **show ip protocols** command or the **debug eigrp** command.

Issue: Make sure that R3 has only a summary of the following routes in its routing table:

- 172.16.55.36/30
- 172.16.55.40/30

- 172.16.55.44/30

Solution:

EIGRP summarization is done on a per-interface basis. It is applied on the interface that you want to send the summary out of, using the command **ip summary-address eigrp**. In this scenario, R5 will have the summary configuration on its interface E0/0. Verify this setup by simply checking the routing table on R3.

R5:

```
interface Ethernet0/0
 ip address 172.16.35.5 255.255.252.0
 ip summary-address eigrp 35 172.16.55.32 255.255.255.240 5
```

R3:

```
R3#sho ip route eigrp | inc 55
D       172.16.55.32/28 [90/156160] via 172.16.35.5, 17:15:05, Ethernet0/1
```

Issue: Advertise the minimal number of prefixes to router R5 but provide full connectivity to R5.

Solution:

What tools do we have to set the gateway of last resort in a router using EIGRP? We have 0.0.0.0/0, but this solution is disallowed. An alternative solution to this configuration requirement is to configure the **ip default-network** command either on router R3 or R5 and advertise a classful prefix to R5. This solution cannot be used because it is restricted in the "Restrictions and Goals" section at the beginning of the scenario. Another solution to consider is local policy routing on router R5, but this is ruled out by the restriction that says that only conventional routing techniques may be used. All of the addresses in the pod are subnets of 172.16.0.0/16, 172.20.0.0/16, or 172.31.0.0/16, so we can provide full reachability by creating a summary to 172.16.0.0/12 on R3 interface E0/1:

```
R3#show run int e0/1
Building configuration...

Current configuration : 196 bytes
!
interface Ethernet0/1
 ip address 172.16.35.3 255.255.252.0
 ip access-group 101 in
 ip bandwidth-percent eigrp 35 25
 ip pim dense-mode
 ip summary-address eigrp 35 172.16.0.0 255.240.0.0
end

R3#
R5#sh ip route | inc /12
D     172.16.0.0/12 [90/307200] via 172.16.35.3, 04:09:03, Ethernet0/0
R5#
```


Issue: Restrict the bandwidth utilization to half of the default value for EIGRP traffic.

Solution:

The **bandwidth-percent** command is configured on the interface and tells EIGRP what percentage of the configured bandwidth it may use. The default is 50 percent, so the percentage in this scenario should be restricted to 25 percent.

R3:

```
interface Ethernet0/0
 ip address 172.16.35.3 255.255.252.0
 ip bandwidth-percent eigrp 35 25
```

R5:

```
interface Ethernet0/0
 ip address 172.16.35.5 255.255.252.0
 ip bandwidth-percent eigrp 35 25
```

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

5. IPv4 Route Redistribution

Before examining the specific issues related to configuring each of the IGPs that are involved in this scenario, it is a good idea to first survey the entire topology and determine how all of the different IGPs will interoperate. Performing such a survey will force you to consider the issues related to route redistribution.

When evaluating a single internetwork topology that contains multiple routing protocols, a good starting point of analysis is to determine whether there is more than one direct or indirect connecting point between two routing protocols. If there is only one connecting point 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 complex. When two or more connecting points exist, you can use them to provide redundancy, as well as load balancing and optimum path selection. However, when two or more connecting points exist, you must also ensure, at the very least, that no routing loops exist and, whenever possible, that no suboptimal paths are selected.

When evaluating this lab internetwork topology, how the routing protocols have been assigned to it and where the specified redistribution points are, you will see that there are at least two possible paths to reach many of the IP addresses assigned in the lab. Of the two possible paths, one path traverses the OSPF domain and the second traverses the RIPv2 domain. Therefore, the only routing protocols providing transit services in this lab are OSPF and RIPv2. The protocol EIGRP is deployed on the edge of the lab topology. RIPv2 provides a transport for BGP peering.

In this scenario, the core protocol is OSPF. It connects RIP and EIGRP. The redistribution points are R1 and R3. There is also a redistribution process on R2, but that only involves connected networks. R1 redistributes between RIP and OSPF. It basically just redistributes RIP into OSPF. In the RIP domain, for reachability, R1 injects a default 0.0.0.0/0 route into the RIP domain and makes R1 the default router for all “points unknown” in RIP.

R3 redistributes between OSPF and EIGRP. There is also a requirement to keep only a minimal number of routes on R5. R3 will take all prefixes from EIGRP and redistribute them into OSPF. However, it will only advertise a 172.16.0.0/12 summary to R5. There are also connected networks for redistribution on R2 and R3. These networks are redistributed into OSPF and will be visible through the OSPF domain as external type 2, or O E2, by default.

Below is a Tool Command Language (Tcl) script that you can use to test universal reachability. To use the script, enter the command **tclsh** in privileged mode, and paste in this script. To kill failing pings, hold down **Ctrl-Shift** and press the **6** key twice. When you are done, enter **tclq** to leave the Tcl mode. This list excludes IP addresses in the CustomerA VPN.

```
tclsh
foreach address {

172.16.14.1
172.16.123.1
172.16.101.1

172.16.25.5
172.16.26.2
172.16.25.1
172.16.25.9
172.20.10.1
172.16.26.100
172.16.123.2
172.16.102.1

172.16.169.1
172.16.168.1
172.16.170.1
172.16.35.3
172.16.123.3
172.16.103.1

172.16.193.1
172.16.192.1
172.16.194.1
172.16.14.4


172.16.35.5
172.16.55.41
172.16.55.45
172.16.55.37
172.16.105.1

172.16.26.6
172.31.1.1
172.16.26.100
172.16.106.1

} {
ping $address
}
```

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

6. MPLS Layer 3 VPN

Issue: Configure Multiprotocol BGP between R1 and R4 using AS number 14. Peer to loopbacks 101 and 104.

Solution:

Internal BGP will be used to transport routing information between the two CustomerA VPN sites. Configure BGP normally, using the Loopback 101 and Loopback 104 IP addresses for peering. Enable the capability to send VPNv4 addresses by activating each neighbor under the VPNv4 address family, as you see here:

R1:

```
router bgp 14
 bgp log-neighbor-changes
 neighbor 172.16.104.1 remote-as 14
 neighbor 172.16.104.1 update-source Loopback101
 !
 address-family vpnv4
  neighbor 172.16.104.1 activate
  neighbor 172.16.104.1 send-community extended
 exit-address-family
```

The command **neighbor 172.16.104.1 send-community extended** is added automatically. Use the command **show bgp vpnv4 unicast all summary** to verify a successful peering.

Configure BGP on R4 similarly to R1.

Issue: Enable the link between R1 and R4 to support VPN labels.

Solution:

To meet this requirement, enter the command **mpls ip** on the R1 interface E0/3 and R4 interface E0/3. A Label Distribution Protocol (LDP) neighbor relationship will form between R1 and R4. The routers will send labels to each other for each network in their routing tables. VPN labels are advertised by BGP, along with each advertised VPN route, and will be affixed to each VPN packet on the link that connects R1 and R4. The command **show mpls ldp neighbor** will verify the adjacency.

```
R1#show mpls ldp neighbor
    Peer LDP Ident: 172.16.104.1:0; Local LDP Ident 172.16.101.1:0
        TCP connection: 172.16.104.1.34088 - 172.16.101.1.646
        State: Oper; Msgs sent/rcvd: 287/269; Downstream
        Up time: 03:47:09
        LDP discovery sources:
          Ethernet0/3, Src IP addr: 172.16.14.4
        Addresses bound to peer LDP Ident:
          172.16.14.4     172.16.104.1    172.16.192.1    172.16.193.1
          172.16.194.1
R1#
```

Issue: Create VPN CustomerA on R1 and R4 using the specified parameters.

Solution:

VPNs are identified by a case-sensitive name on the local router. The 64-bit route distinguisher (RD) is added to the beginning of each route to form the VPNv4 address, and is used to identify its native VPN on other routers. The 64-bit route target (RT) value is set as routes are exported into BGP, and it is sent along with each route as an extended BGP community value. It identifies the originating site, and is used to control which routes are imported into a VRF. Here is the basic configuration on R1 and R4:

```
ip vrf CustomerA
 rd 14:100
 route-target export 14:100
 route-target import 14:100
```

Issue: Place the required interfaces into the CustomerA VRF.

Solution:

On each interface, enter the command **ip vrf forwarding CustomerA**. This removes the interface from the default VRF, as well as the existing IP address. Be sure to re-enter the IP address. These are the results you should see on R1 and R4:

```
R1#show ip vrf
  Name                              Default RD          Interfaces
  CustomerA                         14:100              Et0/0
                                                        Et0/1

R4#show ip vrf
  Name                              Default RD          Interfaces
  CustomerA                         14:100              Et0/1
R4#
```

Issue: Enable EIGRP AS 1 for PE-to-CE, routing. Choose AS numbers to ensure that all learned EIGRP routes are internal.

Solution:

To transport the EIGRP routes between sites, mutually redistribute BGP and EIGRP. The EIGRP AS number from the originating site is advertised by BGP, as is the prefix. EIGRP routes from other sites will be seen as EIGRP external routes if the AS numbers differ across sites, but they will be seen as internal routes if the AS numbers are the same.

Here is the relevant EIGRP configuration. It uses AS number 10, but any AS number would be fine, as long as it is the same at each site.

R4:

```
router eigrp 1
 auto-summary
 !
 address-family ipv4 vrf CustomerA
  redistribute bgp 14 metric 10000 100 255 1 1500
  network 172.16.0.0
  auto-summary
  autonomous-system 10
```


Note there are two AS numbers here: AS 1 for the default VRF and AS 10 for the CustomerA VRF. Be sure to enter an AS number under the VRF for the VPN since it does not inherit the AS number from the default VRF. To get EIGRP routes to the other site, redistribute the PE-to-CE routing protocol into BGP under the **address-family ipv4** command for the VRF, as you see here on R1:

```
router bgp 14
 bgp log-neighbor-changes
 neighbor 172.16.104.1 remote-as 14
 neighbor 172.16.104.1 update-source Loopback101
 !
 address-family ipv4
  neighbor 172.16.104.1 activate
  no auto-summary
  no synchronization
 exit-address-family
 !
 address-family vpnv4
  neighbor 172.16.104.1 activate
  neighbor 172.16.104.1 send-community extended
!
 address-family ipv4 vrf CustomerA
  redistribute eigrp 10
  no synchronization
 exit-address-family
```

Verify the redistribution by checking the local RIB for the VRF. Here is the result on R1:

```
R1#show bgp vpnv4 unicast vrf CustomerA
BGP table version is 13, local router ID is 172.16.101.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 14:100 (default for vrf CustomerA)
*> 172.16.1.0/24    0.0.0.0                  0         32768 ?
*> 172.16.2.0/24    0.0.0.0                  0         32768 ?
*>i172.16.3.0/24    172.16.104.1             0    100      0 ?
*> 172.16.111.0/24  172.16.1.10         156160         32768 ?
*> 172.16.112.0/24  172.16.2.10         156160         32768 ?
*>i172.16.114.0/24  172.16.104.1        156160    100      0 ?
```

Note the locally redistributed connected routes in the VRF, 172.16.1.0/24 and 172.16.2.0/24. You also see the loopback networks on SW1 and SW2, 172.16.111.0/24 and 172.16.112.0/24, which were redistributed from the local EIGRP 10 process. The two IBGP-learned prefixes are the subnets in the CustomerA VRF connected to R4. Note that the BGP next hop for these routes is the IBGP peering address 172.16.104.1. Take a closer look at the attributes of the prefix 172.16.114.0/24:

```
R1#show bgp vpnv4 unicast vrf CustomerA 172.16.114.0
BGP routing table entry for 14:100:172.16.114.0/24, version 13
Paths: (1 available, best #1, table CustomerA)
  Not advertised to any peer
  Refresh Epoch 1
  Local
    172.16.104.1 (metric 1) from 172.16.104.1 (172.16.194.1)
      Origin incomplete, metric 409600, localpref 100, valid, internal, best
      Extended Community: RT:14:100 Cost:pre-bestpath:128:409600
        0x8800:32768:0 0x8801:10:153600 0x8802:65281:256000 0x8803:65281:1500
        0x8806:0:2886758913
      mpls labels in/out nolabel/20
```

```
      mpls labels in/out nolabel/18
      rx pathid: 0, tx pathid: 0x0
R1#
```

Note the full VPNv4 address for this prefix is 14:100:172.16.114.0/24. This would distinguish this prefix from a same subnet originated in another VRF. The highlighted metric value is the metric for this route in R4's routing table. Community 0x8801 shows the originating EIGRP AS number (10) and the scaled delay value for the route. Because the local and far AS numbers match, EIGRP will treat the entire VPN as one EIGRP domain. Here you see this prefix in the routing table of R7:

```
R7#show ip route | beg Gate
Gateway of last resort is not set

      172.16.0.0/16 is variably subnetted, 8 subnets, 2 masks
C        172.16.1.0/24 is directly connected, Ethernet0/0
L        172.16.1.10/32 is directly connected, Ethernet0/0
D        172.16.2.0/24 [90/307200] via 172.16.1.1, 04:29:37, Ethernet0/0
D        172.16.3.0/24 [90/307200] via 172.16.1.1, 04:29:37, Ethernet0/0
C        172.16.111.0/24 is directly connected, Loopback111
L        172.16.111.1/32 is directly connected, Loopback111
D        172.16.112.0/24 [90/435200] via 172.16.1.1, 04:20:05, Ethernet0/0
D        172.16.114.0/24 [90/435200] via 172.16.1.1, 04:29:37, Ethernet0/0
R7#
```

Issue: Verify connectivity within the CustomerA VRF.

Solution:

All of the IP addresses within the CustomerA VPN should be reachable from each other. However, it is possible that reachability between sites will fail, even when routes are fully advertised. VPN traffic is label-switched between R1 and R4. Since they are directly connected, the top label will be null, and only a VPN label will be affixed. Here is the label binding on R4 for 172.16.101.0/24. Note that this subnet is used as the BGP next hop.

```
R4#show mpls ldp bindings 172.16.101.0 24
  tib entry: 172.16.101.0/24, rev 14
        local binding:  tag: 16
        remote binding: tsr: 172.16.101.1:0, tag: imp-null
```

To use this label, which is actually null, R4 must find a precise match in its forwarding table. In other words, communication between the sites will fail if R4 has only a 0.0.0.0/0 route. It would also fail if a 172.16.101.1/32 static route was provided. For this reason, the RIP section permitted two routes to be advertised to R4. Those two routes would need to be the default route and 172.16.101.0/24.

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

7. Router Maintenance

Issue: R3 should return the system date and time to the other routers when they connect to R3 via Telnet to port 13.

Solution:

On R3, enter the global configuration command **service tcp-small-servers**.

Issue: R3 should return time in the EST zone, offset -5 hours.

Solution:

Use the command **clock timezone EST -5** in the global configuration mode on R3.

Test the configuration, here is what the output looks like on R1:

```
R1#telnet 172.16.103.1 13
Trying 172.16.103.1, 13 ... Open
Monday, June 17, 2013 16:10:26-EST

[Connection to 172.16.103.1 closed by foreign host]
R1#
```

Or you can use the **daytime** keyword:

```
R1#telnet 172.16.103.1 daytime
Trying 172.16.103.1, 13 ... Open
Monday, June 17, 2013 16:11:10-EST

[Connection to 172.16.103.1 closed by foreign host]
R1#
```

Issue: Avoid getting to the X28 editor.

Solution:

On R1, disable the service pad with the global configuration command **no service pad**. You can get into this mode by mistakenly typing "x2" when you intend **Ctrl-Shift X 2** to get to another router. Once in this mode, typing **exit** will get you out. But if you don't need the X28 editor, you may as well turn it off. Here is what it looks like:

```
R1#x28

*
*
*
exit

R1#
R1#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R1(config)#no service pad
R1(config)#exit
R1#x28
Translating "x28"
% Unknown command or computer name, or unable to find computer address
```

```
R1#
```

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

8. Security

Issue: Stop sending and receiving ICMP redirects on Ethernet networks.

Solution:

Configure the interface configuration command **no ip redirects**.

Issue: Lower the risk of being an amplifier network for any smurf attacks.

Solution:

This lab is initialized with the **ip directed-broadcast** interface Ethernet0/0 command on R9. Configure the interface configuration command **no ip directed-broadcast**.

Issue: Prevent a quick port scan of UDP ports (Port scan is used by the attacker to find out what services are available on your network).

Solution: When a connection is initiated to a closed UDP port, an "ICMP port unreachable" message is sent to signal that the port is not available. Such messages can be used by an attacker to determine what UDP ports are open on the target system. Configure the interface command **no ip unreachables** to stop sending the messages. This command will disable sending all types of ICMP unreachable messages, limiting amount of information an attacker can collect.

Verification:

```
R9#show run int e0/0
Building configuration...

Current configuration : 103 bytes
!
interface Ethernet0/0
 ip address 172.16.3.4 255.255.255.0
 no ip redirects
 no ip unreachables
end
```

Note: The interface configuration command **no ip directed-broadcasts** is the Cisco IOS Software default. It is not displayed in the **show** output.

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

9. Router QoS

Issue: On R2, limit all traffic originating from the 172.16.26.0/24 network from consuming more than 32,000 b/s of the bandwidth. Also, limit this same classification of traffic from the same network from consuming more than 16,000 b/s of bandwidth for a subset of this traffic that possesses a precedence setting of no higher than 2. Finally, limit this same classification of traffic from the same network from consuming more than 8,000 b/s of bandwidth for an additional subset of this traffic that possesses a precedence setting of 2.

Solution:

This configuration requirement specifies the following three separate limits to bandwidth consumption:

1. Limit the bandwidth consumption of all traffic originating from the 172.16.26.0/24 network to 32 kb/s

2. Limit the bandwidth consumption of all traffic originating from 172.16.26.0/24 and possessing an IP precedence setting of 0, 1 or 2 (no higher than 2) to 16 kb/s.

3. Limit the bandwidth consumption of all traffic originating from 172.16.26.0/24 and possessing an IP precedence setting of 2 to 8 kb/s.

Any QoS requirement specifying the limiting of bandwidth consumption must involve one of two configuration options:

- Traffic shaping
- Traffic policing

Now you must determine which of these options to use. When you carefully read the configuration requirements above, you will notice the phrase “limit this same classification of traffic.” Each sentence that uses this phrase ends with a specification of a “subset” within this “same classification of traffic.” This language is implying a nested MQC policy map, where one policy map sets a limit of bandwidth consumption and a subsequent “nested” MQC policy map further limits the bandwidth consumption of a subset of traffic.

Only traffic policing functions can be nested in this manner. Traffic shapers cannot be nested in this manner. This configuration requirement excludes the use of a nested traffic shaping configuration. Therefore, this is a three-level hierarchical policing configuration requirement.

When configuring hierarchical policers, it is important to remember that the packet classification operation is performed before packet processing. Consider the following diagram:

Three-Level Hierarchical Policer

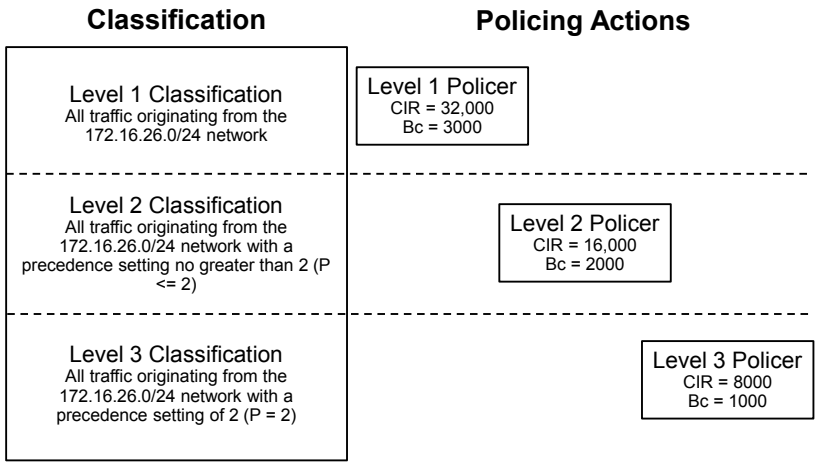

Notice in the diagram above how each preceding level of classification includes each subsequent classification. For example, the Level 3 classification above—matching on packets originating from the 172.16.26.0 network with a precedence of 2—is a subset of the Level 2 classification. This is an important aspect to remember when configuring a hierarchical policer.

Examine how the numbers for the committed information rate (CIR) and committed burst size (Bc) are calculated for each class:

On R2, limit all traffic originating from the 172.16.26.0/24 network from consuming more than 32,000 b/s of the bandwidth.	Level 1 policer	CIR = 32,000 b/s
Also, limit this same classification of traffic from the same network from consuming more than 16,000 b/s of bandwidth for a subset of this traffic that possesses a precedence setting of no higher than 2.	Level 2 policer	CIR = 16,000 b/s
Finally, limit this same classification of traffic from the same network from consuming more than 8000 b/s of bandwidth for an additional subset of this traffic that possesses a precedence setting of 2.	Level 3 policer	CIR = 8000 b/s
Allow for a 750-ms burst for all traffic originating from the 172.16.26.0/24 network and for a 1-second burst of the traffic of other described classes.	Level 1 policer	Bc = 750 ms * 32,000 b/s / 8 = 3000 bytes
	Level 2 policer	Bc = 1 sec * 16,000 b/s / 8 = 2000 bytes
	Level 3 policer	Bc = 1 sec * 8000 b/s / 8 = 1000 bytes

Configuration:

Provided below is the configuration to fulfill the requirements of this section. Hierarchical policing can only be configured with the MQC. As with any MQC configuration, there are three basic sections: two configured in global configuration mode (class maps and policy maps) and one configured in interface configuration mode (service policy).

```
class-map match-all level-1-policer
 match access-group 102
class-map match-all level-2-policer
 match  precedence 0  1  2
 match access-group 102
class-map match-all level-3-policer
 match  precedence 2
 match access-group 102
!
policy-map level-3-policer
 class level-3-policer
   police 8000 1000
policy-map level-2-policer
 class level-2-policer
   police 16000 2000
  service-policy level-3-policer
policy-map level-1-policer
 class level-1-policer
   police 32000 3000
  service-policy level-2-policer
access-list 102 permit ip 172.16.26.0 0.0.0.255 any
```

All of this is applied to the egress R2 Ethernet0/2 interface:

```
interface Ethernet0/2
 service-policy output level-1-policer
```

Note that the traffic policer at the secondary level acts only on packets sent by the policer at the top level. As a result of that, we could have abbreviated the class definitions for Levels 2 and 3 by removing the ACL match (as any packets which passed by top level policer already were selected to match the ACL).The resulting configuration would have looked like this:

```
class-map match-all level-2-policer
 match  precedence 0  1  2
class-map match-all level-3-policer
 match  precedence 2
```

For more information on hierarchical policing, refer Cisco documentation here:

http://www.cisco.com/en/US/docs/ios-xml/ios/qos_plcshp/configuration/15-1mt/qos-plcshp-mod-cli-tlhp.html

Verification:

A very useful command for verifying an MQC configuration is the **show policy-map interface** command. It provides many interesting statistics. Examine the output of this command for the three-level policer configuration. In the output below, notice that it is indented three times, one time for each level of policer.

```
R2#show policy-map inte Ethernet 0/2
```

```
 Ethernet0/2

  Service-policy output: level-1-policer

    Class-map: level-1-policer (match-all)
      0 packets, 0 bytes
      5 minute offered rate 0000 bps, drop rate 0000 bps
      Match: access-group 102
      police:
          cir 32000 bps, bc 3000 bytes
        conformed 0 packets, 0 bytes; actions:
          transmit
        exceeded 0 packets, 0 bytes; actions:
          drop
        conformed 0000 bps, exceeded 0000 bps

      Service-policy : level-2-policer

        Class-map: level-2-policer (match-all)
          0 packets, 0 bytes
          5 minute offered rate 0000 bps, drop rate 0000 bps
          Match:  precedence 0  1  2
          Match: access-group 102
          police:
              cir 16000 bps, bc 2000 bytes
            conformed 0 packets, 0 bytes; actions:
              transmit
            exceeded 0 packets, 0 bytes; actions:
              drop
            conformed 0000 bps, exceeded 0000 bps

          Service-policy : level-3-policer

            Class-map: level-3-policer (match-all)
              0 packets, 0 bytes
              5 minute offered rate 0000 bps, drop rate 0000 bps
              Match:  precedence 2
              Match: access-group 102
              police:
                  cir 8000 bps, bc 1000 bytes
                conformed 0 packets, 0 bytes; actions:
                  transmit
                exceeded 0 packets, 0 bytes; actions:
                  drop
                conformed 0000 bps, exceeded 0000 bps

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

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

      Class-map: class-default (match-any)
        2388 packets, 268840 bytes
        5 minute offered rate 0000 bps, drop rate 0000 bps
        Match: any
R2#
```

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

10. Switch Specialties

Issue: Port 0/3 on SW1 is patched to the visitor room to provide connectivity to different visitors. Allow continuous access to the network for two workstations and allow either workstation to be replaced by another after 5 minutes.

Solution:

The port security feature can help you to solve this issue:

SW1:

```
interface Ethernet0/3
 description Visitor Connectivity
 switchport mode access
 switchport port-security
 switchport port-security maximum 2
 switchport port-security aging time 5
 switchport port-security aging type inactivity
 duplex auto
!
SW1#
```

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

11. Multicast

Issue: Multicast traffic coming from R2 to R1 will be accepted on the Ethernet0/2 interface and will be forwarded to R3 via the tunnel link and to R4 via the link between the Ethernet0/3 interfaces.

The solution is based on introducing another link—a tunnel link between R1 and R3 according to the lab requirements. With the tunnel, R1 will have another interface to forward traffic to R3. The tunnel should be unnumbered, and it should be sourced from and destined to the IP addresses that are assigned to the Ethernet0/2 interfaces in order to conform to the lab requirements.

Multicast Diagram

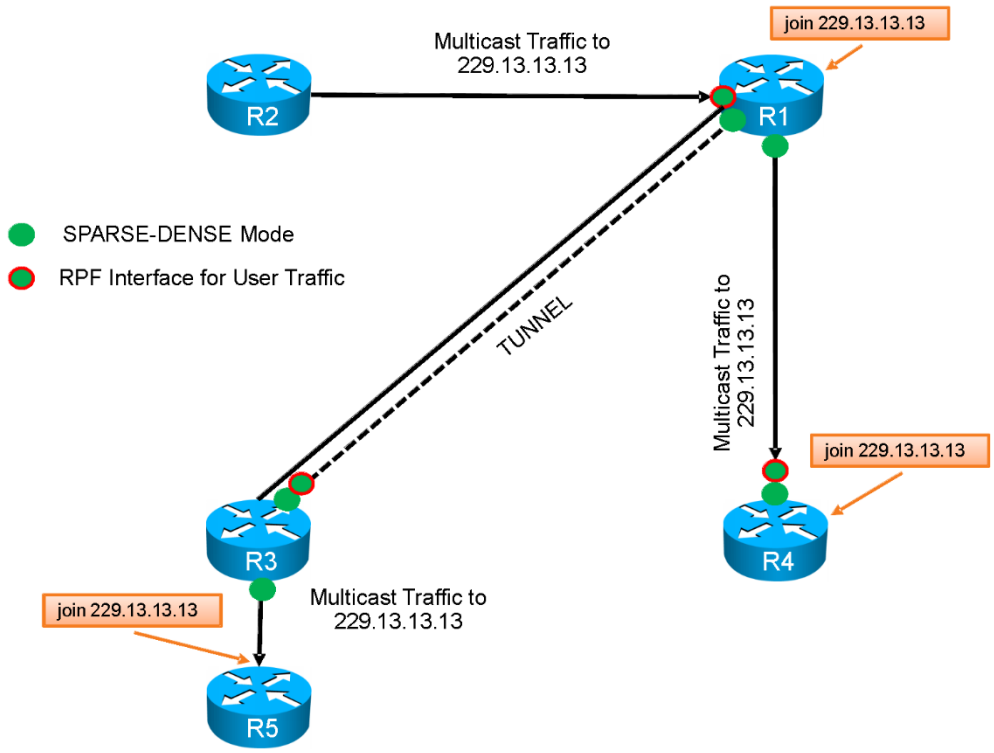

Here is a configuration example on R1:

```
ip multicast-routing
!
interface Loopback101
 ip address 172.16.101.1 255.255.255.0
 ip pim sparse-dense-mode
 ip igmp join-group 229.13.13.13
 ip ospf network point-to-point
!
interface Tunnel10
 description Multicast Link R1-R3
 ip unnumbered Ethernet0/2
 ip pim sparse-dense-mode
 tunnel source Ethernet0/2
 tunnel destination 172.16.123.3
!
interface Ethernet0/2
 ip address 172.16.123.1 255.255.255.0
 ip pim sparse-dense-mode
 ip ospf message-digest-key 1 md5 test
 ip ospf network point-to-multipoint
!
interface Ethernet0/3
 ip address 172.16.14.1 255.255.255.0
 ip pim dense-mode
 mpls ip
```



```
!
```

Here is a configuration example on R3:

```
ip multicast-routing
!
interface Loopback103
 ip address 172.16.103.1 255.255.255.0
 ip pim sparse-dense-mode
 ip igmp join-group 229.13.13.13
 ip ospf network point-to-point
!
interface Tunnel10
 description Multicast Link R1-R3
 ip unnumbered Ethernet0/2
 ip pim sparse-dense-mode
 tunnel source Ethernet0/2
 tunnel destination 172.16.123.1
!
interface Ethernet0/1
 ip address 172.16.35.3 255.255.252.0
 ip access-group 101 in
 ip bandwidth-percent eigrp 35 25
 ip pim sparse-dense-mode
 ip summary-address eigrp 35 172.16.0.0 255.240.0.0
!
interface Ethernet0/2
 ip address 172.16.123.3 255.255.255.0
 ip ospf message-digest-key 1 md5 test
 ip ospf network point-to-multipoint
!
!
```

Note that the Ethernet0/2 interface is not PIM-enabled on R3.

Here is a configuration example on R4:

```
ip multicast-routing
!
interface Loopback192
 ip address 172.16.192.1 255.255.255.0
 ip pim sparse-dense-mode
 ip igmp join-group 229.13.13.13
!
interface Ethernet0/3
 ip address 172.16.14.4 255.255.255.0
 ip accounting precedence input
 ip pim sparse-dense-mode
 mpls ip
!
!
```

Here is a configuration example on R5:

```
interface Ethernet0/0
 ip address 172.16.35.5 255.255.252.0
 ip bandwidth-percent eigrp 35 25
 ip summary-address eigrp 35 172.16.55.32 255.255.255.240
 ip igmp join-group 229.13.13.13
!
```

Here is an example of the PIM neighbor relationships on R1:

```
R1#show ip pim neighbor
PIM Neighbor Table
Mode: B - Bidir Capable, DR - Designated Router, N - Default DR Priority,
      P - Proxy Capable, S - State Refresh Capable, G - GenID Capable
Neighbor          Interface                Uptime/Expires    Ver   DR
Address                                                            Prio/Mode
172.16.123.3      Tunnel10                 08:02:41/00:01:43 v2    1 / S P G
172.16.14.4       Ethernet0/3              08:02:49/00:01:43 v2    1 / DR S P G
R1#
```

The tunnel interface is not the RPF interface for the source of the multicast traffic 172.16.123.2 on R3, so the traffic coming across the tunnel fails the RPF check and is dropped.

```
R3#show ip rpf 172.16.123.2
 failed, no route exists
R3#

R2# ping 229.13.13.13 source 172.16.123.2
Type escape sequence to abort.
Sending 1, 100-byte ICMP Echos to 229.13.13.13, timeout is 2 seconds:
Packet sent with a source address of 172.16.123.2

Reply to request 0 from 172.16.101.1, 1 ms
Reply to request 0 from 172.16.192.1, 1 ms
R2#
```

Note that R2 is receiving ping replies from R1 and R4, but not from R3 and R5

To fix this problem, enter the following static mroute entry on R3:

```
ip mroute 0.0.0.0 0.0.0.0 tunnel10

R2#ping 229.13.13.13 source 172.16.123.2
Type escape sequence to abort.
Sending 1, 100-byte ICMP Echos to 229.13.13.13, timeout is 2 seconds:
Packet sent with a source address of 172.16.123.2

Reply to request 0 from 172.16.101.1, 1 ms
Reply to request 0 from 172.16.35.5, 1 ms
Reply to request 0 from 172.16.103.1, 1 ms
Reply to request 0 from 172.16.192.1, 1 ms
R2#

R1#show ip mroute 229.13.13.13 172.16.123.2
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector
Outgoing interface flags: H - Hardware switched, A - Assert winner
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(172.16.123.2, 229.13.13.13), 00:02:38/00:00:21, flags: LT
  Incoming interface: Ethernet0/2, RPF nbr 0.0.0.0
```



```
  Outgoing interface list:
    Loopback101, Forward/Sparse-Dense, 00:02:38/stopped
    Ethernet0/3, Forward/Dense, 00:02:38/stopped
    Tunnel10, Forward/Sparse-Dense, 00:02:38/stopped

R1#

R3#show ip mroute 229.13.13.13 172.16.123.2
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector
Outgoing interface flags: H - Hardware switched, A - Assert winner
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(172.16.123.2, 229.13.13.13), 00:03:47/00:02:42, flags: LT
  Incoming interface: Tunnel10, RPF nbr 0.0.0.0, Mroute
  Outgoing interface list:
    Loopback103, Forward/Sparse-Dense, 00:03:47/stopped
    Ethernet0/1, Forward/Sparse-Dense, 00:03:47/stopped

R3#
```

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

12. Gateway Redundancy

Issue: On VLAN 40, configure a first-hop router selection solution that uses UDP port 3222 for end systems. The technique selection must represent multiple routers via a single IP address. The existence of the multiple routers must be completely transparent to the end systems.

Solution:

When reading this configuration requirement, three possible configuration options come to mind: HSRP, VRRP, and GLBP.

All three options allow you to represent multiple routers via a single logical IP address. All three options perform this function in a manner that is transparent to the end user.

Now, the question becomes which option must be chosen. The options are narrowed down to only one choice by the requirement that the solution use UDP port 3222. This choice would be GLBP.

Issue: Make R2 the router that coordinates the allocation of MAC addresses in ARP responses. Make R6 its backup router.

Solution:

This requirement is directing you to configure R2 as the GLBP active virtual gateway (AVG). By default, all GLBP routers are potential AVGs. A GLBP AVG is elected in a manner similar to how the active HSRP router is elected: by the highest priority value. Like HSRP, the GLBP priority value has a default priority value of 100. This can be manually adjusted in a manner identical to HSRP. More information on the specific configuration commands will be supplied in a later section. Before we review the configuration requirements needed to fulfill this specific section, spend a few moments to better understand the basic operation of GLBP. Review the diagram below:

Gateway Redundancy Diagram

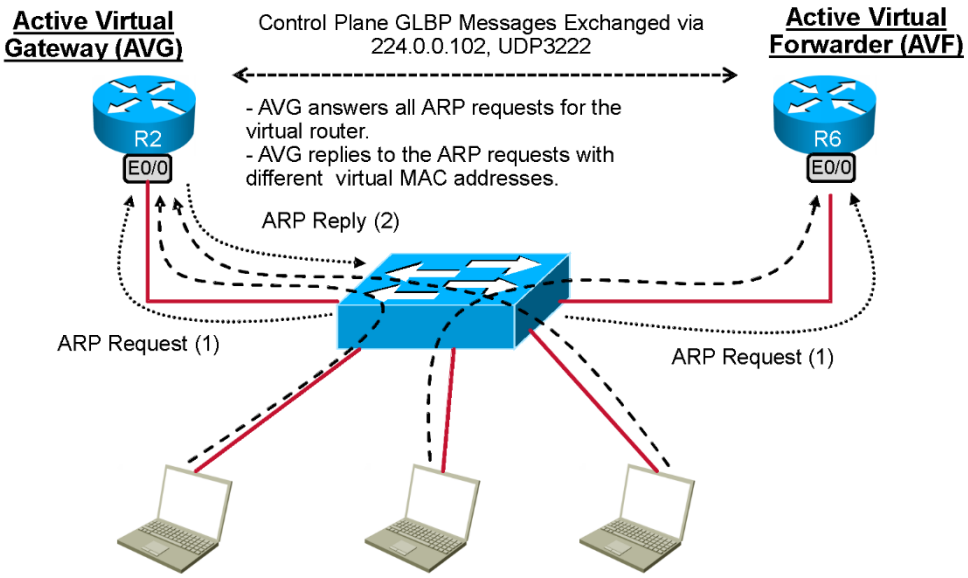

As the diagram reflects, end system workstations broadcast out ARP requests, and the GLBP AVG answers these ARP requests with a set of MAC addresses. Unlike HSRP, where the active HSRP router responds to ARP requests with only its own MAC address, the GLBP AVG responds with other participating routers' MAC addresses. This is how the load-balancing function is performed. In the diagram, the AVG sometimes responds to ARP requests with a locally assigned MAC address. However, it also responds with the MAC addresses of other GLBP participating routers. In the diagram, notice how router R6 is designated as an active virtual forwarder (AVF). As the diagram reflects, R2 (the AVG) will at times respond to ARP requests with a MAC address assigned to an AVF (in this case, R6). Again, this is how GLBP

performs its load-balancing function. GLBP can use different load-balancing techniques, such as round robin and weighted load-balancing. The default load-balancing technique is round robin.

Issue: To fulfill the requirements of this section, limit the participation to routers R2 and R6. Fulfill this requirement by not using any global configuration commands.

Solution:

The language of this configuration requirement forces you to consider a GLBP authentication option. There are two GLBP authentication options, the **key-string** option and the **key-chain** option.

The **key-string** option only requires a single interface configuration command for activation. The **key-string** option requires no global configuration commands. However, the **key-chain** option requires both a global configuration command and an interface configuration command. Because the requirement explicitly forbids any global configuration requirements to fulfill this task, only the **key-string** option can be used.

Configuration:

Start with the most minimal GLBP configuration. As you see, the minimal GLBP configuration requires only one command on each router participating in a GLBP group:

R2:

```
interface Ethernet0/0
 ip address 172.16.26.2 255.255.255.0
 glbp 1 ip 172.16.26.100
end
```

R6:

```
interface Ethernet0/0.40
 ip address 172.16.26.6 255.255.255.0
 glbp 1 ip 172.16.26.100
end
```

That is all that is needed to perform a basic GLBP configuration.

To manually fix the election of the AVG, enter a GLBP configuration command that is very similar to the HSRP configuration command used to fix an HSRP active router election:

```
glbp 1 priority XXX
```

The default GLBP priority value is 100. A higher GLBP priority value is more preferred when electing a GLBP AVG. Since this configuration requirement directed you to make R2 the AVG, you should have configured R2 with a GLBP priority value of greater than 100. The value of 110 is used in this answer key.

While it was not specified, a GLBP pre-emption feature was also configured. The GLBP pre-emption feature operates in the same manner as the HSRP pre-emption feature. It is a recommended general practice to use this feature unless you are explicitly instructed to do otherwise.

Configuring the GLBP Authentication Feature

To fulfill the authentication requirement for this section, enter in the following interface configuration command on all GLBP participating interfaces on R2 and R6:

```
glbp 1 authentication md5 key-string test
```

Verification:

To check the basic configuration of GLBP, enter the following **show** command. It provides much useful information about the basic operation of GLBP.

```
R6#sh glbp
Ethernet0/0.40 - Group 1
  State is Standby
    1 state change, last state change 00:00:50
  Virtual IP address is 172.16.26.100
  Hello time 3 sec, hold time 10 sec
    Next hello sent in 0.396 secs
  Redirect time 600 sec, forwarder time-out 14400 sec
  Preemption disabled
  Active is 172.16.26.2, priority 100 (expires in 9.588 sec)
  Standby is local
  Priority 100 (default)
  Weighting 100 (default 100), thresholds: lower 1, upper 100
  Load balancing: round-robin
  Group members:
    0004.9ada.b0c0 (172.16.26.6) local
    00d0.5895.c8e1 (172.16.26.2)
  There are 2 forwarders (1 active)
  Forwarder 1
    State is Listen
    MAC address is 0007.b400.0101 (learnt)
    Owner ID is 00d0.5895.c8e1                          <= This is R2
    Time to live: 14399.588 sec (maximum 14400 sec)
    Preemption enabled, min delay 30 sec
    Active is 172.16.26.2 (primary), weighting 100 (expires in 9.588 sec)
  Forwarder 2
    State is Active
      1 state change, last state change 00:00:57
    MAC address is 0007.b400.0102 (default)
    Owner ID is 0004.9ada.b0c0                          <= This is R6
    Preemption enabled, min delay 30 sec
    Active is local, weighting 100
```

To obtain a precise understanding of the operation of GLBP, it is best to view not only GLBP **show** commands but router and workstation ARP tables, as well. Here are some sample ARP tables from the GLBP participating routers:

```
R2#sh arp
Protocol  Address          Age (min)  Hardware Addr   Type   Interface
Internet  172.16.26.2             -   aabb.cc00.0200  ARPA   Ethernet0/0
Internet  172.16.26.6            78   aabb.cc00.0600  ARPA   Ethernet0/0
Internet  172.16.26.100            -   0007.b400.0102  ARPA   Ethernet0/0
Internet  172.16.123.1            94   aabb.cc00.0120  ARPA   Ethernet0/2
Internet  172.16.123.2             -   aabb.cc00.0220  ARPA   Ethernet0/2
Internet  172.16.123.3            88   aabb.cc00.0320  ARPA   Ethernet0/2
R2#

R6#sh arp
Protocol  Address          Age (min)  Hardware Addr   Type   Interface
Internet  172.16.26.6              -   0001.4219.2a40  ARPA   Ethernet0/0
Internet  172.16.26.2             25   0004.9adb.7181  ARPA   Ethernet0/0
Internet  172.16.26.100            -   0007.b400.0102  ARPA   Ethernet0/0
```


To test the operation of GLBP, configure the following on one VLAN using three routers:

1. Configure one router to act as a client workstation.

2. Configure at least two routers to act as GLBP routers.

3. Enable **debug arp** and **debug ip icmp** on each of the participating routers.

4. Clear the ARP cache on the router that is acting as the client workstation.

5. Ping an unreachable remote address.

6. Check the ARP table on the workstation.

7. View the **debug arp** message. ARP replies for the default gateway address always come from the AVG.

8. However, the ICMP messages are generated in a round robin.

9. Finally, check the "ARP replies sent" counter in the following **show** command:

```
R2#sh glbp | b Forwarder
Forwarder 1
    State is Listen
      4 state changes, last state change 00:01:19
    MAC address is 0007.b400.0101 (learnt)
    Owner ID is aabb.cc00.0600
    Redirection enabled, 598.496 sec remaining (maximum 600 sec)
    Time to live: 14398.496 sec (maximum 14400 sec)
    Preemption enabled, min delay 30 sec
    Active is 172.16.26.6 (primary), weighting 100 (expires in 10.336 sec)
    Arp replies sent: 9
  Forwarder 2
    State is Active
      1 state change, last state change 00:01:47
    MAC address is 0007.b400.0102 (default)
    Owner ID is aabb.cc00.0200
    Redirection enabled
    Preemption enabled, min delay 30 sec
    Active is local, weighting 100
    Arp replies sent: 9
R2#
```

Once you have the basic operation of GLBP working, you can configure GLBP authentication. As the configuration section reflected, you can do so using a single interface configuration command on each GLBP participating router. Here is a **show** command that you can use to validate the operation of GLBP authentication:

```
R6#sh glbp | i Authentication
  Authentication MD5, key-string

R6:
!
interface Ethernet0/0
 ip address 172.16.26.6 255.255.255.0
 glbp 1 ip 172.16.26.100
 glbp 1 preempt
 glbp 1 authentication md5 key-string test
```

!

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

