360RSW05-WKB-CA04

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 4 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 4 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 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 151.10.0.0/16.
- *Points will be deducted from multiple sections for failing to assign correct IPv4 addresses.*
- Do not use any static routes.
- Advertise loopback interfaces with their original masks.
- Network 0.0.0.0/0 should not appear in any routing table (**show ip route**).
- Do not use the **default-information originate**, **ip default-gateway**, and **ip default-network** commands.
- Do not introduce any new IP addresses.
- All IP addresses involved in this scenario must be reachable, unless explicitly specified otherwise.

- Unless explicitly specified otherwise, addresses and networks that are advertised in the BGP section need to be reachable by all BGP routers but do not have to be reachable by routers that only use IGP.
- 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 diagram belong to network 151.10.0.0/16.

All IP addresses in this lab belong to the 151.10.0.0/16 address space, except for prefixes that are explicitly specified as being part of a different IP space.

Do not use any static routes.

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

Advertise loopback interfaces with their original masks.

The original mask is the mask that is configured on the loopback interface. By default, Open Shortest Path First (OSPF) treats loopback interfaces as host routes and advertises them as /32 prefixes. You need to provide a solution to represent the original mask of addresses assigned to loopback interfaces. One possible option is to change the OSPF network type or summarizations.

Network 0.0.0.0/0 should not appear in any routing table (show ip route**), except on SW2.**

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 default-information originate**,** ip default-gateway**, or** ip default-network **commands.**

These commands can be used to solve reachability issues by setting the gateway of last resort. They generate a 0.0.0.0/0 route. You cannot use them in this scenario.

All IP addresses involved in this scenario must be reachable.

This goal is a key goal to observe. It requires that all of your IGPs be 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.

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

This statement relaxes the requirement that all IP addresses must be reachable. The BGP prefixes need only be reachable among the routers that are specified in the BGP section. It is acceptable if they are 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 known via BGP.

Use conventional routing algorithms only, unless specified otherwise.

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. They 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

Like any switch configuration, you must address the following basic configuration requirements:

- Set the VLAN Trunking Protocol (VTP) mode.
- Configure the VLANs and the VLAN names.
- Configure the trunk ports.
- Statically assign the ports of the switches to the VLANs.

Note	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. Some of the Cisco Catalyst 3560 Switch configuration commands are not available on the virtual instances of the switches.

Use the "VLANs" table, the "Switch-to-Router Connections" table, and the "Switch-to-Switch Connections" table for reference.

Make sure that the VLAN names are spelled correctly. VLAN names are case-sensitive.

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

See the following diagram for the VLAN layout.

VLAN Distribution

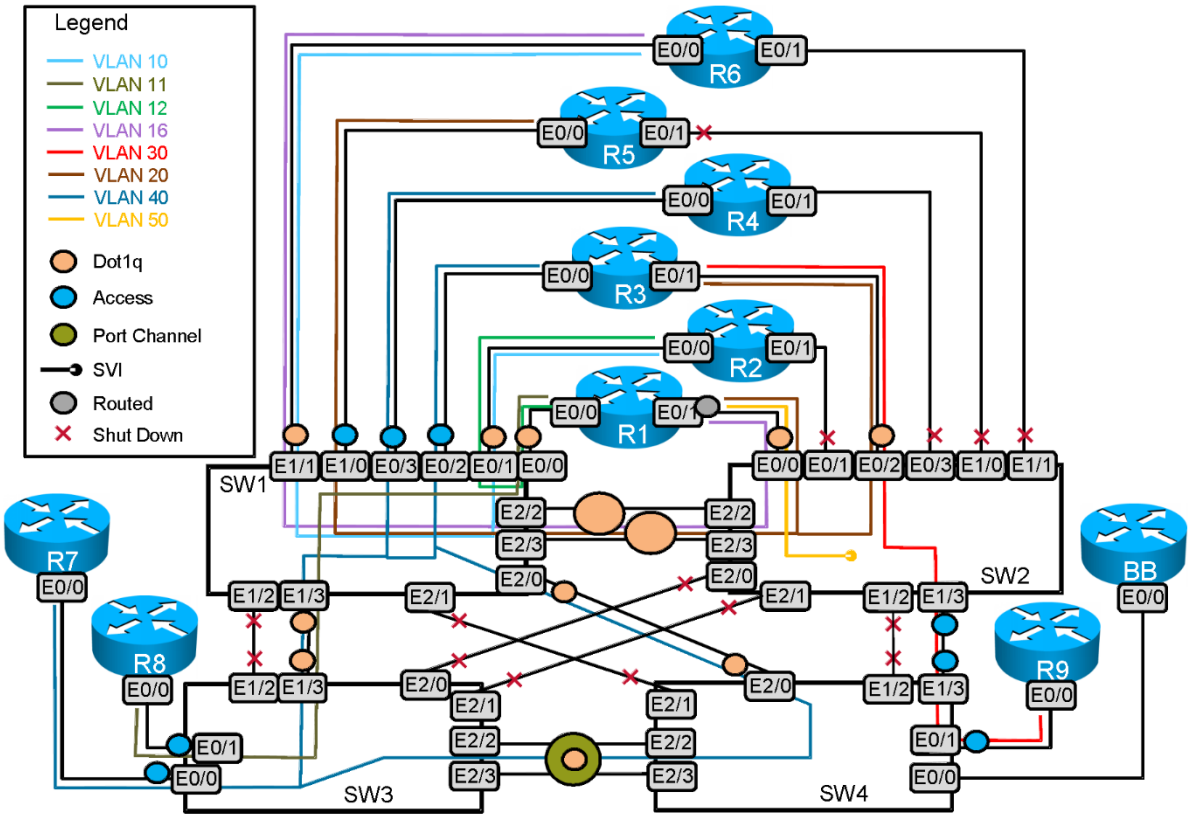

Issue: Do not allow switches to advertise VLANs.

Solution:

If you are told not to advertise VLANs in a scenario, you will need to configure VTP transparent or off mode on all of your switches. Switches that are configured in VTP transparent mode do not advertise any VLANs that are created on them.

To verify VTP status on the switch, issue the **show vtp status** command:

```
VTP Version                     : 2
Configuration Revision          : 0
Maximum VLANs supported locally : 1005
Number of existing VLANs        : 9
VTP Operating Mode              : Transparent
VTP Domain Name                 :
VTP Pruning Mode                : Disabled
VTP V2 Mode                     : Disabled
VTP Traps Generation            : Disabled
MD5 digest                      : 0xC3 0xB7 0x9D 0xD2 0xF6 0xF3 0xDE 0x44
Configuration last modified by 0.0.0.0 at 0-0-00 00:00:00
```

Issue: SW4 should load-balance between interfaces 2/2 and 2/3.

Solution:

Configure interfaces 2/2 and 2/3 of SW3 and SW4 in a port channel.

Issue: Make sure that SW3 is the root for VLAN 40. Hellos should be sent every 3 seconds on VLAN 40 and be considered valid for 19 seconds. The Hello forward delay should be set to 14 seconds.

Solution:

Apply the following configuration on SW3:

```
SW3#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
SW3(config)#spanning-tree vlan 40 root primary
SW3(config)#spanning-tree vlan 40 hello-time 3
SW3(config)#spanning-tree vlan 40 max-age 19
SW3(config)#end
SW3#
```

Verify the VLAN 40 spanning tree configuration on SW3:

```
SW3#show spanning-tree vlan 40

VLAN0040
  Spanning tree enabled protocol ieee
  Root ID     Priority    24616
              Address     aabb.cc00.0900
              This bridge is the root
              Hello Time   3 sec  Max Age 19 sec  Forward Delay 14 sec

  Bridge ID   Priority    24616  (priority 24576 sys-id-ext 40)
              Address     aabb.cc00.0900
              Hello Time   3 sec  Max Age 19 sec  Forward Delay 14 sec
              Aging Time  300 sec

Interface           Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- --------------------------------
-
Et0/0               Desg FWD 100       128.1    Shr
Et1/3               Desg FWD 100       128.8    Shr
Po1                 Desg FWD 56        128.65   Shr


SW3#
```

Issue: When sending frames for VLAN 40, make sure SW1 prefers interface 2/0 to interface 1/3.

Solution:

SW1 receives bridge protocol data units (BPDUs) from the following:

- SW3 on the E1/3 interface , with cost = 0

```
SW1#show spanning-tree vlan 40 interface ethernet1/3 detail
 Port 8 (Ethernet1/3) of VLAN0040 is root forwarding
   Port path cost 100, Port priority 128, Port Identifier 128.8.
   Designated root has priority 24616, address aabb.cc00.0900
   Designated bridge has priority 24616, address aabb.cc00.0900
   Designated port id is 128.8, designated path cost 0
   Timers: message age 1, forward delay 0, hold 0
   Number of transitions to forwarding state: 1
```



```
   Link type is shared by default
   BPDU: sent 385, received 496
SW1#
```

- SW4 on the E0/2 interface, with cost = 56

```
SW1#show spanning-tree vlan 40 interface ethernet2/0 detail
 Port 9 (Ethernet2/0) of VLAN0040 is alternate blocking
   Port path cost 100, Port priority 128, Port Identifier 128.9.
   Designated root has priority 24616, address aabb.cc00.0900
   Designated bridge has priority 32808, address aabb.cc00.0a00
   Designated port id is 128.9, designated path cost 56
   Timers: message age 4, forward delay 0, hold 0
   Number of transitions to forwarding state: 1
   Link type is shared by default
   BPDU: sent 835, received 52
SW1#
```

SW1 then adds 100 to the received cost.

Therefore, SW1 prefers the path through SW3 (cost 100+0 = 100) to the path through SW4 (cost 100+56=156).

To ensure that SW1 prefers the path through SW4, you must increase the cost of interface 1/3 on SW1 to any value greater than 156.

```
SW1#show spanning-tree vlan 40

VLAN0040
  Spanning tree enabled protocol ieee
  Root ID     Priority    24616
              Address     aabb.cc00.0900
              Cost        156
              Port        9 (Ethernet2/0)
              Hello Time   3 sec  Max Age 19 sec  Forward Delay 14 sec

  Bridge ID   Priority    32808  (priority 32768 sys-id-ext 40)
              Address     aabb.cc00.0700
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
              Aging Time  14  sec

Interface           Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- --------------------------------
-
Et0/2               Desg FWD 100       128.3    Shr
Et0/3               Desg FWD 100       128.4    Shr
Et1/3               Altn BLK 157       128.8    Shr
Et2/0               Root LIS 100       128.9    Shr


SW1#
```

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. IPv4 OSPF

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

Issue: Choose the OSPF network type that will form a collection of point-to-point adjacencies on the subnet 151.10.124.0/24 and will generate host entries for the Ethernet0/2 interfaces.

Solution:

The 151.10.124.0/24 subnet is an Ethernet network. If you configure one of the two following OSPF network types—point-to-multipoint or point-to-multipoint nonbroadcast—you will generate a host route for each interface that is connected to the subnet.

Configure the point-to-multipoint OSPF network type by entering the **ip ospf network point-to-multipoint** command in the interfaces of the 151.10.124.0/24 subnet.

Issue: Choose the OSPF nonbroadcast network type on the serial link between R3 and R4.

Solution:

The default OSPF network type for a serial interface is point-to-point. Configure the serial link between R3 and R4 as an OSPF nonbroadcast network type and provide the OSPF neighbor configuration.

Issue: Assign loopback 151.10.20.1/24 to area 20 on R3.

Solution:

This task sets the stage for the very last configuration requirement in the OSPF section. See the last configuration requirement in this section for more details.

Issue: Assign loopback 151.10.44.129/25 to area 44 on R4.

Solution:

The “Restrictions and Goals” section instructs learners to advertise loopback interfaces with their original masks. When loopback interfaces are assigned to an OSPF area, they are advertised as host routes by default. To change this behavior, configure the loopback interface as a point-to-point OSPF network type. With this configuration, the IP address that is assigned to the loopback interface on R4 will be advertised with its native prefix.

Configure the loopback with the point-to-point OSPF network type by issuing the **ip ospf network point-to-point** command.

Issue: On R3, add loopback 151.10.10.1/24 in OSPF without using a network statement, interface configuration OSPF statements, or the **redistribute connected** command.

Solution:

Since you can't use the network statement or the **redistribute connected** command, redistribute the prefix into OSPF via another routing protocol that is running on R3. When you read ahead in this scenario, you will notice that Routing Information Protocol (RIP) is also configured on R3. All 151.10.0.0 prefixes, including 151.10.10.0./24, are part of the RIP process. Therefore, you can redistribute this prefix into OSPF via RIP.

Issue: Make R7 the preferred candidate for the designated router (DR), and R4 the preferred candidate for the backup designated router (BDR).

Solution:

There are three routers on the 151.10.34.0/24 subnet. To make R7 the preferred DR candidate, set its OSPF priority to the highest value. To make R4 the preferred BDR candidate, set its OSPF priority to the highest value minus 1. R3 is a potential candidate for both DR and BDR, but the configuration task does not specify how it should be ranked in the DR or BDR election process. So set its OSPF priority to *0* so that it can never override R7 and R4.

Configure priority level 0 on R3 by issuing the **ip ospf priority 0** command.
Configure priority level 2 on R7 by issuing the **ip ospf priority 2** command.
R4 will retain its default OSPF priority level 1 on interface Et0/0.

On R3, issue the **show ip ospf neighbor** command:

```
Neighbor ID     Pri   State           Dead Time   Address         Interface
151.10.100.4      1   FULL/BDR        00:01:38    151.10.34.4     Ethernet0/0
151.10.100.7      1   FULL/DR         00:01:36    151.10.34.7     Ethernet0/0
```

Issue: Make sure that R4 can reach network 151.10.20.0 using the shortest possible OSPF path.

Solution:

Configure a virtual link over VLAN 40 to make it the preferred path to the 151.10.20.0 prefix from R4. If you do not do this, R4 will always prefer the slower serial link path over the faster Ethernet path because OSPF will always select an area 0 path over a non-area 0 path. By configuring the virtual link on VLAN 40, area 0 is extended over this VLAN. Once the virtual link is configured over VLAN 40, R4 will have two area 0 paths to the 151.10.10.0 and 151.10.20.0 prefixes. R4 will then select the Ethernet path over the Serial path, since the Ethernet path has a lower OSPF cost.

On R4, issue the **area 10 virtual-link 151.10.100.3** command.
On R3, issue the **area 10 virtual-link 151.10.100.4** command.

On R4, issue the **show ip route ospf** command and verify that 151.10.20.0/24 is reachable over the Ethernet link:

```
R4#show ip route ospf | inc 20.0
O IA     151.10.20.0/24 [110/11] via 151.10.34.3, 00:42:33, Ethernet0/0
R4#
```

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. IPv4 EIGRP

Issue: Prevent Enhanced Interior Gateway Routing Protocol (EIGRP) queries from being sent to SW2. SW2 should be able to ping the rest of the network.

Solution:

The only router that can send EIGRP queries to SW2 is R1. Configure SW2 with the **eigrp stub** command. This will prevent R1 from sending EIGRP queries to SW2.

Issue: Prevent R1 from learning 140.10.1.0/25 and 140.10.2.0/24 routes via EIGRP. Deny the 140.10.2.0/24 route by matching the IP prefix, and deny the 140.10.1.0/25 route by matching the source protocol. Permit all other routes.

Solution:

You can use the EIGRP Route Map Support feature to accomplish this task. This feature allows routes to be filtered using route maps, and enables matching EIGRP attributes such as metric or source protocol.

R1:

```
router eigrp 100
 distribute-list route-map EIGRP-filter in Ethernet0/1.16
!
ip prefix-list EIGRP-filter seq 5 permit 140.10.2.0/24
!
route-map EIGRP-filter deny 10
 match source-protocol connected
!
route-map EIGRP-filter deny 20
 match ip address prefix-list EIGRP-filter
!
route-map EIGRP-filter permit 30
```

Note that R1 can potentially learn two external EIGRP routes: 140.10.2.0/24 and 151.10.100.2/32. The first one can only be redistributed into EIGRP as connected, while the second one can be redistributed either as connected or as RIP. Redistribute 151.10.100.2/32 into EIGRP from RIP on R2 to ensure that the above route map does not filter it out (which would have broken the requirements of the VPN section).

Shut down the Ethernet0/2 interface on R1 and examine the routing table.

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. IPv4 RIP

Issue: Configure RIP version 2 (RIPv2) only between R1, R2, R3, and R5.

Solution:

R1 and R3 will be running RIPv2. R5 will speak only RIP.

Issue: Make RIP send updates three times more often than they are sent by default.

Solution:

Set the RIP update timer from the default of 30 seconds to 10 seconds in the **router-rip** configuration process.

Issue the **timers basic 10 60 60 80** command as part of the router RIP configuration.

Issue the **show ip protocol** command and observe the following information about RIP:

```
Routing Protocol is "rip"
  Sending updates every 10 seconds, next due in 6 seconds
  Invalid after 60 seconds, hold down 60, flushed after 80
```

Note that Cisco IOS Software Release 12.4 adds the **ip rip advertise 10** command to all interfaces that are covered by the RIP network statement.

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. Redistribution

In this scenario, OSPF plays the role of the core routing protocol. EIGRP and RIP are edge protocols and do not provide transit services. The connection between RIP and OSPF is redundant (dual point), with redistribution on R1 and R3. Administrative distance was used on R1, R2 and R3 so that RIP domain routes are forwarded using internal paths; this is a best practice and is not required by this lab. Connection between EIGRP and OSPF is nonredundant (single-point) with redistribution on R1. VPNs and BGP create interesting complications for redistribution. We will review these in the relevant sections.

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

Redistribution Table

Redist. Point	Into RIP		Into OSPF		Into EIGRP	
	Permit	Deny	Permit	Deny	Permit	Deny

R1	**All routes from OSPF All routes from EIGRP**		**All routes from EIGRP RIP native routes only from RIP**		**All routes from OSPF RIP native routes only from RIP**	
R2			**Connected 151.10.100.2/32 140.10.2.0/24 EIGRP 140.10.1.0/25**		**RIP 151.10.100.2/32**	
R3	**All routes from OSPF**		**RIP native routes only from RIP**			
R4			**Connected 151.10.100.4/32**			
R6					**Connected 140.10.1.0/25**	
SW3			**Connected 151.10.100.7/32**			

You can use the following Tool Command Language (Tcl) script to test universal reachability. To use the script, enter the command **tclsh** in privileged mode, and paste in the script. To kill failing pings, hold down **Ctrl-Shift** and press the **6** key twice. When you are finished, enter **tclquit** to leave Tcl mode.

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

```
tclsh
foreach ip {
151.10.124.1
151.10.12.1
198.1.1.1
151.10.22.1
151.10.135.1
151.10.124.2
151.10.26.2
151.10.12.2
151.10.100.2
151.10.43.3
151.10.34.3
151.10.135.3
151.10.10.1
151.10.20.1
151.10.100.3
151.10.43.4
151.10.124.4
151.10.34.4
151.10.44.129
151.10.100.4
151.10.135.5
151.10.100.5
151.10.26.6
198.1.1.6
151.10.66.1
151.10.100.6
140.10.1.1
151.10.22.20
151.10.34.7
```



```
151.10.100.7
} {ping $ip}
```

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. BGP

Issue: Use the synchronization method in this section.

Solution:

When you are instructed to use the synchronization method in BGP, this means you must make sure that all IBGP learned prefixes have a corresponding prefix listed in the local forwarding table if that IBGP prefix is to be advertised to another BGP speaker. This is commonly fulfilled by performing one-way redistribution on an External Border Gateway Protocol (EBGP) router into the IGP of the autonomous system (AS). Therefore, in order for synchronization to be relevant, you must have an IBGP speaker.

When reading the configuration tasks in this BGP configuration requirement, notice that a total of four autonomous systems are involved: AS 20, AS 134, AS 52, and AS 50. Only AS 134 has more than one BGP speaker. Therefore, only AS 134 has the potential for possessing an IBGP speaker. In other words, the issue of synchronization is only relevant for AS 134.

The default setting for synchronization depends on the version of Cisco IOS Software being used. Issue the **synchronization** command for router BGP in AS 134.

Issue: Add network 140.10.1.0/25 in the R2 BGP routing process.

Solution:

This 140.10.1.0/25 prefix is at the center of the BGP configuration challenge for this scenario. It is simple enough that it could be originated in BGP. However, you must remember that this prefix was originally advertised by R6 via EIGRP. As the task in the "Redistribution" section mandates, you redistributed this prefix into OSPF from EIGRP on R2, which made R2 the Autonomous System Boundary Router (ASBR) for the 140.10.1.0/25 prefix. An add-on to the rule of synchronization is that if the underlying IGP is OSPF for any specific prefix, then the router ID (RID) of the ASBR must equal the RID of the IBGP speaker that advertises the same prefix.

Originating the BGP update on R2 means that there will probably be an RID mismatch between the OSPF ASBR and the advertising IBGP speaker for the 140.10.1.0/25 prefix in AS 134.

Issue: Advertise the network 140.10.2.0/24 on R2. Both networks 140.10.1.0/25 and 140.10.2.0/24 should be added to BGP on R2 without using redistribution commands. Mark network 140.10.1.0/25 with community 20:3 and network 140.10.2.0/24 with community 20:5. Both R1 and R2 will see the prefixes marked with communities.

Solution:

This requirement leads to a quality of service (QOS) configuration. To introduce the prefixes into the R2 configuration, use the **network** command. The **network** command as part of the BGP

routing process inserts a prefix into the BGP table, provided that the prefix is present in the local routing table. It does not matter how the prefix enters the local routing table; one of the prefixes is known as Connected, while the other is learned via EIGRP. Attach a route map to the network commands to provide the desired community marking. Also configure R2 to send community information to R1. (By default, community information is not sent to BGP neighbors.)

R2

```
router bgp 20
 network 140.10.1.0 mask 255.255.255.128 route-map BGP-community-20:3
 network 140.10.2.0 mask 255.255.255.0 route-map BGP-community-20:5
 neighbor 151.10.124.1 remote-as 134
!
route-map BGP-community-20:5 permit 10
 set community 20:5
!
route-map BGP-community-20:3 permit 10
 set community 20:3
```

Issue: Make sure that networks 140.10.1.0/25 and 140.10.2.0/24 reside in the BGP table on R5.

Solution:

R5 will never receive the BGP prefixes unless *all* of the synchronization issues are fulfilled on R3. In other words, R3 will never advertise the 140.10.1.0/25 prefix to R5 unless *all* of the synchronization issues are fulfilled. It is not enough to have 140.10.1.0/25 in the IGP table of R3.

IPv4 BGP Diagram

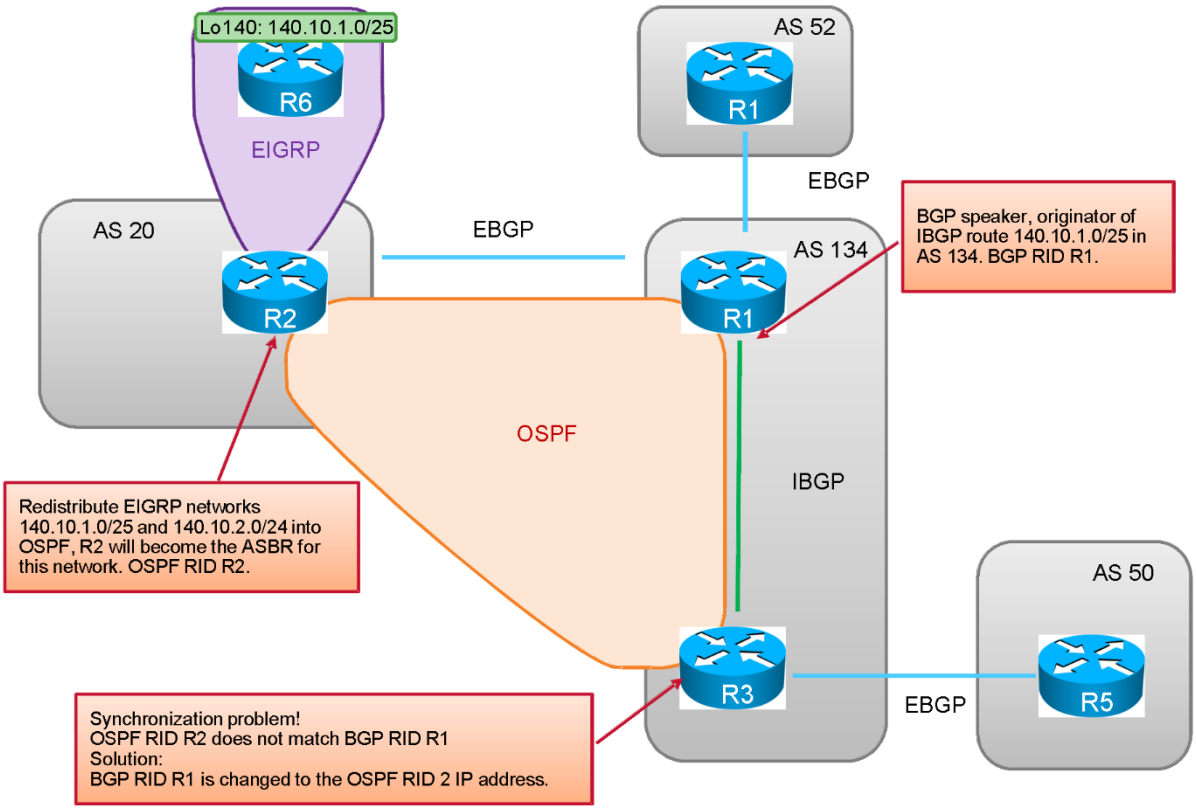

As mentioned in the previous section, the RID of the OSPF ASBR and the RID of the advertising IBGP speaker must match to fulfill the synchronization requirements when OSPF is the underlying IGP. This can be accomplished by switching the BGP router IDs on routers R1 and R2, provided that both the RIDs of OSPF and BGP were selected dynamically.

Note that both OSPF and BGP select their RIDs in an identical manner. If you switch the BGP RIDs on routers R1 and R2, R1, the advertising IBGP speaker in AS 134 will possess the BGP RID that matches the RID of the OSPF ASBR—router R2. Once this is done, all synchronization requirements are fulfilled on router R3. As a result, R3 will advertise the 140.10.1.0/25 prefix to R5 and the configuration requirements for this section are fulfilled.

Issue: The BGP connection between R1 and R2 should always be initiated from R2. Apply the configuration on R1.

Solution:

The usual configuration of BGP results in a router that originates and accepts TCP sessions for BGP peering. Use the **neighbor** *ip-address* **transport connection-mode** command to modify the default behavior.

R1:

```
router bgp 134
```

```
 neighbor 151.10.124.2 transport connection-mode passive
```

Use the **show ip bgp neighbors** command to verify configuration. Note that if the default configuration is used, the command output does not include information on how the BGP session will be opened:

```
R1#show ip bgp neighbors 151.10.124.2 | i (^BGP)|(TCP)
BGP neighbor is 151.10.124.2,  remote AS 20, external link
  TCP session must be opened passively
TCP Semaphore      0xF090EB34  FREE
R1#
```

The **show tcp** command can be used to examine active TCP sessions:

```
R1#show tcp brief
TCB       Local Address           Foreign Address        (state)
6598E998  151.10.135.1.52299      151.10.135.3.179       ESTAB
6598FA70  151.10.22.1.38140       151.10.22.20.179       ESTAB
6598C8F0  151.10.124.1.179        151.10.124.2.20563     ESTAB
R1#
```

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. MPLS Layer 3 VPN

Issue: Create VPN 14 on R1 and R3.

Solution:

Import and export route targets are not specified, so you will have to choose values that meet the overall requirements. A simple and common approach is to use the route distinguisher value for both route targets. Here is an example configuration:

```
ip vrf VPN14
 rd 134:14
 route-target export 134:14
 route-target import 134:14
```

You then associate the required interfaces with the VRF using the command **ip vrf forwarding VPN14**.

Verify with the command **show ip vrf**:

```
R1#show ip vrf
  Name                              Default RD          Interfaces
  VPN14                             134:14              Lo14
                                                        Et0/0.11

R3#show ip vrf
  Name                              Default RD          Interfaces
  VPN14                             134:14              Lo14
                                                        Et0/1.30
R3#
```


Note that adding the interface to the VRF removes it from the default VRF; it will no longer be listed when you enter **show ip route connected**. Remember to add back the original IP address.

Issue: Configure OSPF for provider-edge-to-customer-edge routing.

Solution:

The challenge here is to remember to use an appropriate OSPF process ID and the **vrf** keyword. When this step is complete, you should see a routing table like this:

```
R1#show ip route vrf VPN14

Routing Table: VPN14

Gateway of last resort is not set

     151.10.0.0/24 is subnetted, 5 subnets
O       151.10.14.0 [110/101] via 151.10.1.11, 00:01:26, Ethernet0/0.11
O       151.10.4.0 [110/102] via 151.10.1.11, 00:01:26, Ethernet0/0.11
C       151.10.1.0 is directly connected, Ethernet0/0.11
O       151.10.114.0 [110/102] via 151.10.1.11, 00:01:26, Ethernet0/0.11
O       151.10.111.0 [110/2] via 151.10.1.11, 00:01:26, Ethernet0/0.11
```

R8 and R9 are connected by a link that is intended to be a backup path between the sites, in case the MPLS core is unavailable. At this point, all of the VPN 14 IP addresses are reachable using this backup link. Here is a trace from R1 to 151.10.4.3 on R3:

```
R1#trace vrf VPN14 151.10.4.3

Type escape sequence to abort.
Tracing the route to 151.10.4.3

  1 151.10.1.11 4 msec 0 msec 0 msec
  2 151.10.14.4 4 msec 0 msec 0 msec
  3 151.10.4.3 4 msec *  0 msec
```

Issue: Configure BGP to carry VPN routes between R1 and R3 and enable the transmission of VPN 14 traffic across VLAN 20. Mutually redistribute BGP and OSPF.

Solution:

R1 and R3 created an IBGP peering in AS 134 in the earlier section. Enter the **neighbor activate** command under the VPNv4 address family on both R1 and R3. The **neighbor send-community extended** command will be added automatically. Verify the peering as you see here:

```
R1#show bgp vpnv4 unicast all summary | begin Neighbor
Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
151.10.135.3    4   134     770     772       36    0    0 00:07:51        0
```

On both R1 and R3, mutually redistribute between BGP and OSPF, as you see here:

```
router ospf 14 vrf VPN14
```

```
 log-adjacency-changes
 redistribute bgp 134 subnets
 network 151.10.1.0 0.0.0.255 area 14

router bgp 134
 [output removed for brevity]
 !
address-family vpnv4
  neighbor 151.10.135.3 activate
  neighbor 151.10.135.3 send-community extended
 exit-address-family
 !
 address-family ipv4 vrf VPN14
  redistribute ospf 14 vrf VPN14
  no synchronization
 exit-address-family
```

When this step is complete, you should see BGP and routing tables like this on R1:

```
R1#show bgp vpnv4 unicast vrf VPN14
BGP table version is 40, local router ID is 151.10.100.2
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: 134:14 (default for vrf VPN14)
* i151.10.1.0/24    151.10.135.3           102    100      0 ?
*>                  0.0.0.0                  0         32768 ?
*> 151.10.4.0/24    151.10.1.11            102         32768 ?
* i                  151.10.135.3             0    100      0 ?
* i151.10.14.0/24   151.10.135.3           101    100      0 ?
*>                  151.10.1.11            101         32768 ?
* i151.10.111.0/24  151.10.135.3           102    100      0 ?
*>                  151.10.1.11              2         32768 ?
*> 151.10.114.0/24  151.10.1.11            102         32768 ?
* i                  151.10.135.3             2    100      0 ?


R1#show ip route vrf VPN14 | begin Gateway
Gateway of last resort is not set

     151.10.0.0/24 is subnetted, 5 subnets
O       151.10.14.0 [110/101] via 151.10.1.11, 00:01:55, Ethernet0/0.11
O       151.10.4.0 [110/102] via 151.10.1.11, 00:01:55, Ethernet0/0.11
C       151.10.1.0 is directly connected, Ethernet0/0.11
O       151.10.114.0 [110/102] via 151.10.1.11, 00:01:55, Ethernet0/0.11
O       151.10.111.0 [110/2] via 151.10.1.11, 00:01:55, Ethernet0/0.11
```

Note that the BGP paths across the MPLS core are not placed into the local forwarding table, even though the path across the MPLS core is faster than the backup route.

VPN traffic between R1 and R3, across VLAN 20, must carry a VPN label. Enter the command **mpls ip** on both interfaces to enable MPLS label switching between R1 and R3. Verify with the command **show mpls ldp neighbor**.

Issue: Optimize site-to-site traffic.

Solution:

Why is the slow backup path preferred over the fast MPLS path? The answer becomes apparent if we temporarily shut the backup link and check the routing table on R8:


```
R8#sh ip route | begin Gateway
Gateway of last resort is not set

     151.10.0.0/24 is subnetted, 4 subnets
O IA    151.10.4.0 [110/2] via 151.10.1.1, 00:00:52, Vlan11
C       151.10.1.0 is directly connected, Vlan11
O IA    151.10.114.0 [110/3] via 151.10.1.1, 00:00:52, Vlan11
C       151.10.111.0 is directly connected, Loopback111
```

Note that the paths across the MPLS core are OSPF interarea routes, and have a much lower OSPF cost. However OSPF prefers intra-area routes over interarea routes; when both paths are available, OSPF prefers the backup intra-area paths even though interarea paths have a higher cost.

To address this problem, configure an OSPF sham link between R1 and R3. The sham link creates an adjacency in the common OSPF area, so the routes learned across the MPLS core will be intra-area, rather than interarea, and can be compared based on cost.

OSPF sham link end points must be host addresses in the VPN. To avoid recursive routing they are redistributed into BGP as connected routes, rather than included in the OSPF process. We created loopback 14 on R1 with IP address 14.14.14.1/32 and loopback 14 on R3 with IP address 14.14.14.3/32, added them into VPN14, and redistributed connected under the BGP ipv4 address family for the VRF. Finally, we entered the sham link commands under the VPN OSPF process on R1 and R3:

```
R1#sh run | begin ospf 14
router ospf 14 vrf VPN14
 log-adjacency-changes
 area 14 sham-link 14.14.14.1 14.14.14.3
 redistribute bgp 134 subnets
 network 151.10.1.0 0.0.0.255 area 14

R3#show run | begin ospf 14
router ospf 14 vrf VPN14
 log-adjacency-changes
 area 14 sham-link 14.14.14.3 14.14.14.1
 redistribute bgp 134 subnets
 network 151.10.4.0 0.0.0.255 area 14
```

Verify the sham link operation with the commands **show ip ospf sham-link** and **show ip ospf neighbor**.

With the sham link in place, there is an OSPF intra-area path across the MPLS core:

```
R1#show ip route vrf VPN14 | begin Gateway
Gateway of last resort is not set

     151.10.0.0/24 is subnetted, 5 subnets
O       151.10.14.0 [110/101] via 151.10.1.11, 00:09:20, Ethernet0/0.11
O       151.10.4.0 [110/2] via 151.10.135.3, 00:09:20
C       151.10.1.0 is directly connected, Ethernet0/0.11
O       151.10.114.0 [110/3] via 151.10.135.3, 00:09:20
O       151.10.111.0 [110/2] via 151.10.1.11, 00:09:20, Ethernet0/0.11
     14.0.0.0/32 is subnetted, 2 subnets
B       14.14.14.3 [200/0] via 151.10.135.3, 00:46:03
C       14.14.14.1 is directly connected, Loopback14
```

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. IPv6 Routing

Issue: Configure link-local and site-local IPv6 addresses.

Solution:

The link-local address of an interface may need to be referenced by another router on the same link. Example of where such a reference may be necessary include routing protocol neighbor statements (such as OSPFv3, IPv6 EIGRP, and IPv6 BGP)

Manual configuration of link-local addresses is not required for such a reference. The automatically assigned link-local address of the router can be examined using the **show ipv6 interface** command, and can be referenced in a neighbor router configuration. In some lab scenarios, there is a requirement to manually configure the link-local address.

Manual configuration of a link-local address is also a good practice with production networks in which the link-local address must be referenced in another router configuration.

Assign IPv6 addresses to all interfaces. At this point, you should be able to ping IPv6 addresses from each router and reach the interfaces of all the other routers that are sharing the same link. You thus verify that connectivity on a link basis has been achieved.

Enable IPv6 unicast routing on all routers at this time. Unlike IPv4 unicast routing, IPv6 unicast routing is not enabled by default:

```
ipv6 unicast-routing
```

Issue: Configure OSPF area 0 between R1 and R2. Use the point-to-point OSPF network type on the Ethernet links. Ensure that the Ethernet0/2 link is preferred for traffic forwarding.

Solution:

Change the default OSPFv3 broadcast network type to the OSPFv3 point-to-point network type by using the **ipv6 ospf network point-to-point** command on the Ethernet interfaces of R1 and R2. Use the **ipv6 ospf cost** command to prefer a path for forwarding.

Issue: On the VLAN 12 link, declare the neighbor inactive when hello packets are not received for a period of 20 seconds. Do not change the default number of hellos that can be missed before the neighbor is declared inactive.

Solution:

The OSPFv3 neighbor activity and inactivity timers depend on the network type that you configure on the link. This lab requires the point-to-point network time to be configured on the VLAN 12 link. The default OSPFv3 timers for the point-to-point network type are shown in the following example on R1:

```
R1#show ipv6 ospf interface Ethernet0/0.12
Ethernet0/0.12 is up, line protocol is up
  Link Local Address FE80:151:10:12::1, Interface ID 16
  Area 0, Process ID 1, Instance ID 0, Router ID 198.1.1.1
  Network Type POINT_TO_POINT, Cost: 100
  Transmit Delay is 1 sec, State POINT_TO_POINT
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:05
  Graceful restart helper support enabled
  Index 1/1/1, flood queue length 0
  Next 0x0(0)/0x0(0)/0x0(0)
  Last flood scan length is 5, maximum is 5
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 151.10.100.2
  Suppress hello for 0 neighbor(s)
R1#
```

The default dead interval—the time period for which hello packets must not be seen before neighbors declare the router inactive—equals four times the hello interval, which is the interval between hello packets that the Cisco IOS Software sends on the interface. To fulfill the requirement, you need to maintain the ratio (four hellos per dead interval). You can adjust the dead interval proportionally by modifying the hello interval using the **ipv6 ospf hello-interval 5** command.

Issue: Assign the R2 loopback 102 interface to OSPF area 1. All other OSPF routers should see this prefix summarized to /64.

Solution:

As with IPv4, IPv6 OSPF ABR can summarize area routes to a larger prefix.

R2

```
interface Loopback102
 ipv6 address 2001:151:10:100::2/128
 ipv6 ospf 1 area 1
!
ipv6 router ospf 1
 area 1 range 2001:151:10:100::/64
```

Issue: Assuming that the OSPF network topology for IPv6 is stable, suppress unnecessary flooding of link state advertisements (LSAs) on the connection between the Ethernet0/2 interfaces of R1 and R2.

Solution:

Like IPv4 implementation, IPv6 OSPF supports flood reduction that is configured with the **ipv6 ospf flood-reduction** command.

R1

```
interface Ethernet0/2

 ipv6 ospf 1 area 0
 ipv6 ospf network point-to-point
 ipv6 ospf flood-reduction
!
```

R2

```
interface Ethernet0/2

 ipv6 ospf 1 area 0
 ipv6 ospf network point-to-point
 ipv6 ospf flood-reduction
```

Issue: Configure IPv6 EIGRP AS 1 between R1 and R6, and between R2 and R6.

Solution:

R1

```
interface Ethernet0/1.16
 ipv6 address 2001:198:1:1::1/64
 ipv6 address FE80:198:1:1::1 link-local
 ipv6 eigrp 1
!
ipv6 router eigrp 1
```

R2

```
interface Ethernet0/0.10
 ipv6 address 2001:151:10:26::2/64
ipv6 eigrp 1
!
ipv6 router eigrp 1
```

R6:

```
interface Ethernet0/0.10
 ipv6 address 2001:151:10:26::6/64
ipv6 eigrp 1
-
!
interface Ethernet0/0.16
 ipv6 address 2001:198:1:1::6/64
 ipv6 address FE80:198:1:1::6 link-local
ipv6 eigrp 1
!
ipv6 router eigrp 1
```

Issue: Provide reachability for R1 network 2001:151:10:22::/64 without adding it to any routing protocols.

IPv6 reachability between the R1 interface on VLAN 50 and R2 loopback 102 should be configured in a redundant way as follows: the primary path should be via the Ethernet0/2 link, and the secondary path should be via the VLAN 12 link. The tertiary path should be via the R6 link.

Solution:

There are two redistribution requirements here. The first requirement is in reference to the R1 VLAN 50 subnet. The second requirement addresses the need to provide a redundant path between two networks, which requires that the two prefixes in question are present in both the OSPFv3 and IPv6 EIGRP routing tables. Since the primary path is over OSPF, the default administrative distance does not need to be modified.

Note that the R2 loopback prefix is summarized for OSPF. If you advertise it in IPv6 EIGRP without summarization, it results in a more specific (not summarized) prefix being preferred, based on the longest-match rule. Administrative distances are compared only between prefixes of the same length, while prefixes of different lengths arrive at the routing table independently; the choice as to which one to use is governed solely by the longest-match rule. To avoid this issue, summarize the prefix in IPv6 EIGRP, too.

R1

```
ipv6 router eigrp 1
 redistribute connected metric 1 1 1 1 1
!
ipv6 router ospf 1
 redistribute connected
!
```

R2

```
interface Ethernet0/0.10
 ipv6 summary-address eigrp 1 2001:151:10:100::/64
!
ipv6 router eigrp 1
 redistribute connected metric 1 1 1 1 1
!
ipv6 router ospf 1
 area 1 range 2001:151:10:100::/64
 redistribute connected
!
```

Issue: Configure R1 to use the IPv6 DNS server that is located at address 2001::1. Do not enable DNS-based name-to-address translation.

Allow R1 to perform Telnet to loopback 102 of R2 using the name “R2-v6” for the IPv6 protocol, and the name “R2” for the IPv4 protocol.

Solution:

Cisco IOS Software supports DNS over both IPv4 and IPv6 protocols. Configure the name server as usual using the IPv6 address. The local name entries are similar to their IPv4 equivalent, but with a minor syntax difference:

R1:

```
no ip domain lookup
ip host R2 151.10.100.2
ip name-server 2001::1
ipv6 host R2-v6 2001:151:10:100::2

R1#telnet R2-v6
Translating "R2-v6"
Trying 2001:151:10:100::2 ... Open
```

```
-----------------------------------------------
Cisco 360 R&S Exercise Workbook
Product, POD location:  cierswbv5-ce-lab04-sc, SJ
Device:                 R2
-----------------------------------------------

R2#exit

[Connection to R2-v6 closed by foreign host]
R1#
```

Issue: On R1, modify the defaults for sending Internet Control Message Protocol (ICMP) error messages. Make the token bucket twice as large and add tokens twice as often.

Solution:

This quote is taken from *Cisco IOS IPv6 Configuration Guide, Release 12.4*, in the configuration section called "Configuring IPv6 ICMP Rate Limiting":

"In Cisco IOS Release 12.2(8)T or later releases, the IPv6 ICMP rate limiting feature implements a token bucket algorithm for limiting the rate at which IPv6 ICMP error messages are sent out on the network. The initial implementation of IPv6 ICMP rate limiting defined a fixed interval between error messages, but some applications, such as traceroute, often require replies to a group of requests sent in rapid succession. The fixed interval between error messages is not flexible enough to work with applications such as traceroute and can cause the application to fail. Implementing a token bucket scheme allows a number of tokens—representing the ability to send one error message each—to be stored in a virtual bucket. The maximum number of tokens allowed in the bucket can be specified, and for every error message to be sent, one token is removed from the bucket. If a series of error messages is generated, error messages can be sent until the bucket is empty. When the bucket is empty of tokens, IPv6 ICMP error messages are not sent until a new token is placed in the bucket. The token bucket algorithm does not increase the average rate-limiting time interval, and it is more flexible than the fixed time interval scheme."

To modify token bucket parameters, use the command **ipv6 icmp error-interval.** Refer to the command reference for the defaults. Use the **show ipv6 traffic** command to examine ICMP statistics.

```
ipv6 icmp error-interval 50 20
```

After you have completed the configuration of all routing tasks for this scenario, the following **tclsh** script can be used to verify universal reachability:

```
tclsh
foreach ip {
2001:151:10:124::1
2001:151:10:12::1
2001:198:1:1::1
2001:151:10:22::1
2001:151:10:124::2
2001:151:10:26::2
2001:151:10:12::2
2001:151:10:100::2
2001:151:10:26::6
2001:198:1:1::6
2001:151:10:66::1
} {ping $ip}
```


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. QoS

Issue: Traffic sent by SW2 to BGP networks in AS 20 should be marked on R1 with IP precedence 3 for prefixes marked with BGP community 20:3, and with IP precedence 5 for prefixes that are marked with BGP community 20:5.

Solution:

The BGP section has led to this requirement by advertising prefixes that are marked with communities. To use these communities for packet marking, do the following:

Ensure that the router receives prefixes marked with communities.

Enable Cisco Express Forwarding switching on the router using the **ip cef** command..

On R1, create the community lists and the route map to verify the communities and set the IP precedence. Reference the route map in the **router-bgp** configuration.

```
ip bgp-community new-format
ip community-list 3 permit 20:3
ip community-list 5 permit 20:5
!
route-map BGP-QOS permit 10
 match community 3
 set ip precedence flash
!
route-map BGP-QOS permit 20
 match community 5
 set ip precedence critical
!
router bgp 20
 table-map BGP-QOS
```

At this point, you should be able to verify the Cisco Express Forwarding table using the **show ip cef** command.

Configure the ingress interface to perform classification:

```
interface Ethernet0/1.50
 bgp-policy destination ip-prec-map
```

The command can classify based on either the source or destination IP address (this would be the IP address that is used to look up the Cisco Express Forwarding table), and can set either the IP precedence or the QoS group.

Verify that BGP routes are received with expected communities:

```
R1#show ip bgp 140.10.1.0
BGP routing table entry for 140.10.1.0/25, version 2
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Advertised to non peer-group peers:
  151.10.135.3
  20
```

```
    151.10.124.2 from 151.10.124.2 (2.2.2.2)
      Origin IGP, metric 20, localpref 100, valid, external, best
      Community: 20:3

R1#show ip bgp 140.10.2.0
BGP routing table entry for 140.10.2.0/24, version 3
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Advertised to non peer-group peers:
  151.10.135.3
  20
    151.10.124.2 from 151.10.124.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, valid, external, best
      Community: 20:5
```

Verify that the precedence information is in the Cisco Express Forwarding table:

```
R1#show ip cef 140.10.1.0 detail
140.10.1.0/25, epoch 0, flags rib only nolabel, rib defined all labels
  QOS: Precedence flash (3)
  recursive via 151.10.124.2
    nexthop 151.10.124.2 Ethernet0/2
R1#

R1#show ip cef 140.10.2.0 detail
140.10.2.0/24, epoch 0, flags rib only nolabel, rib defined all labels
  QOS: Precedence critical (5)
  recursive via 151.10.124.2
    nexthop 151.10.124.2 Ethernet0/2
R1#
```

Create test traffic and verify classification. You may want to consider temporarily creating an access control list (ACL) on the target router to permit all IP traffic, but separately count the traffic of the marking that has been established:

```
access-list 199 permit ip any any precedence 3 log
access-list 199 permit ip any any precedence 5 log
access-list 199 permit ip any any
```

Attach the ACL to the inbound interface on the target router and confirm the marking. Do not forget to remove ACL when testing is completed.

Issue: The R1 Ethernet0/2 interface leads to another QoS domain. Modify the QoS marking on the packets that are leaving the interface as follows:

QoS Mappings

Original QoS Value	Class	New QoS Value
ip precedence 5	1	ip dscp ef
ip precedence 2	2	ip dscp af31
ip precedence 3	3	ip dscp af33
Anything else	4	ip dscp 0

Solution:

This task requires the marking of outbound traffic.

Define the classes of traffic:

```
class-map match-all QOS-1
 match ip precedence 5
```



```
class-map match-all QOS-2
 match ip precedence 2
class-map match-all QOS-3
 match ip precedence 3
```

Create a service policy to provide the marking that is required:

```
policy-map QOS
 class QOS-1
  set dscp ef
 class QOS-2
  set dscp af31
 class QOS-3
  set dscp af33
class class-default
  set dscp default
```

Attach the service policy to the Ethernet0/2 interface:

```
interface Ethernet0/2
 service-policy output QOS
```

Issue: Allocate 30 percent of bandwidth to Class 1 and provide priority treatment for Class 1 traffic. For bandwidth that is not allocated to priority class, use 20 percent for Class 2 and 30 percent for Class 3. Provide preferential treatment of interactive traffic within Class 4.

Solution:

This task requires the addition of class-based weighted fair queuing (CBWFQ) to the previously created configuration.

First examine the bandwidth specification. The bandwidth for the priority queue is specified in percent and not as an absolute value (such as kb/s).

In addition, the bandwidth for the remaining classes is specified as a percentage of bandwidth not allocated to the priority class versus to the full bandwidth available.

Cisco IOS Software introduced the Low Latency Queuing with Priority Percentage Support feature, which enables the fulfillment of both configuration requirements. The percentage-based specification is helpful when a single policy map is used in conjunction with a number of interfaces that have different bandwidth parameters. Note that although it is possible to provide the same allocation by manually recalculating the percentage into kb/s and then using calculated values to specify the bandwidth, such a configuration would not meet the requirements of the task (although they would be the same functionally).

Now look at the Class 4 requirement. The default class within the CBWFQ structure allows for the use of flow-based fair queuing, which provides preferential treatment to interactive flows. It is not possible to use fair queuing with any other class.

```
policy-map QOS
 class QOS-1
  priority percent 30
  set dscp ef
 class QOS-2
  bandwidth remaining percent 20
  set dscp af31
 class QOS-3
```

```
  bandwidth remaining percent 30
  set dscp af33
class class-default
  fair-queue
  set dscp default
```

Create some test traffic and observe the CBWFQ operation. Use the **show policy-map** command. Also consider the tracking ACL method described above to confirm re-marking of the traffic.

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. System Administration

Issue: The network administrator has observed that R5 has a high CPU utilization. Implement a configuration so that R5 CPU utilization is examined every 5 seconds and information is sent to 140.10.2.10 server when the CPU utilization exceeds 50 percent, when the CPU utilization for activities other than interrupt processing exceeds 30 percent, and when CPU utilization returns to below 10 percent.

Solution:

Use the CPU Thresholding Notification feature to accomplish this task. This feature significantly simplifies configuration of CPU utilization monitoring.

Utilization can be monitored in total or separately for interrupt and process-based activity.

R5

```
snmp-server enable traps cpu threshold
snmp-server host 140.10.2.10 public
process cpu threshold type total rising 50 interval 5 falling 10 interval 5
process cpu threshold type process rising 30 interval 5 falling 10 interval 5
```

The following messages will be present in the local log when thresholds are crossed:

```
%SYS-1-CPURISINGTHRESHOLD: Threshold: Total CPU Utilization
%SYS-1-CPUFALLINGTHRESHOLD: Threshold: Total CPU Utilization

%SYS-1-CPURISINGTHRESHOLD: Threshold: Process CPU Utilization
%SYS-1-CPUFALLINGTHRESHOLD: Threshold: Process CPU Utilization
```

Simple Network Management Protocol (SNMP) traps will also be sent to the server.

Issue: Configure R5 to accept Secure Shell (SSH) connections. Do not change the domain name configuration on R5. Only SSH version 2 (SSHv2) connections will be accepted. Use the key named R5.cisco.com and the minimal allowable modulus.

Configure R2 to establish SSHv2 sessions to the loopback 0 interface of R5 when **connect-to-r5** is entered.

Solution:

To enable SSH, generate the crypto key using the **crypto key generate** command.

To specify the SSH version, use the **ip ssh version** command.

By default, the router acting as the SSH server will accept connections of either SSHv1 or SSHv2. To restrict the router so that it accepts only connections of a specific version, use the **ip ssh version** command.

Usually, routers acting as the SSH server need to be configured with a host name and domain name. To configure the SSH server on R5 without specifying a domain name, generate the crypto usage key with the label specified and configure SSH to use this key with the **ip ssh rsa keypair-name** command.

Note that although a crypto key can be generated with any modulus from 512 to 2048 bits, for SSHv2, the modulus size must be at least 768 bits.

```
crypto key generate rsa usage-keys label R5.cisco.com modulus 768
ip ssh rsa keypair-name R5.cisco.com
ip ssh version 2
```

The **crypto key generate** command is not reflected in the configuration because key generation is a one-time event. The generated keys are stored in a secure area of NVRAM and can be examined using the **show crypto key mypubkey rsa** command:

```
R5#show crypto key mypubkey rsa
% Key pair was generated at: 13:13:27 PST May 28 2013
Key name: R5.cisco.com
Key type: RSA KEYS
 Storage Device: not specified
 Usage: General Purpose Key
 Key is not exportable.
 Key Data:
  307C300D 06092A86 4886F70D 01010105 00036B00 30680261 00BF33A9 B56D9500
  E951C1AB EAC8C638 07602AD6 90690905 A29D3379 1686E308 068F7572 C5BE0FD6
  51209667 250557D2 D39A02E1 C4A631AB 1DA88883 F807A7B6 FF9E8055 41D2630E
  4E585BC8 10956380 AAB8251D 7AB12AAA 6E1C4DE4 F58159E0 91020301 0001
WARNING: List changed. Show command aborted.
R5#
```

SSHv2 functionality was introduced for both client and server code. Use the **ssh** command from exec mode to initiate SSH sessions.

Configure the **username cisco password cisco** and **login local** commands on R5.

The requirement specifies that the alias to the **ssh** command be created on R2:

R2

```
alias exec connect-to-r5 ssh -l cisco -v 2 151.10.100.5
```

Ensure that the SSH version and user name are specified in the command. The **ssh** command, unlike Telnet, does not support interactive entry of the user name.

```
R2#connect-to-r5
Password:

------------------------------------------------
Cisco 360 R&S Exercise Workbook
Product, POD location:  cierswbv5-ce-lab04-sc, SJ
Device:                 R5
------------------------------------------------
```

```
R5#exit

[Connection to 151.10.100.5 closed by foreign host]
R2#
```

Issue: Configure R5 to improve the configuration management of the router by enabling faster collection of running configuration-file information.

Solution:

This task requires the configuration of the Configuration Generation Performance Enhancement feature in Cisco IOS Software. When enabled, it improves speed-configuration information to be built for a running configuration by caching configuration information in memory.

R5
```
parser config cache interface
```

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. Address Administration

Issue: From R4, open an HTTP connection to the IP address 151.10.10.1 and get the router prompt from R5. Apply the relevant configuration on R3.

Solution:

This is an **ip nat inside source static** NAT configuration requirement that not only maps an IP address but also maps a TCP port.

R3
```
interface Serial1/0
 ip nat outside
interface Ethernet0/0
 ip nat outside
interface Ethernet0/1.20
 ip nat inside
!
ip nat inside source static tcp 151.10.135.5 23 151.10.10.1 80 extendable
```

In addition, examine the route redistribution configuration on R1 and R3. R5 will recognize the source IP address of the incoming TCP session as one of the R4 addresses. For example:

```
R5#who
    Line       User       Host(s)              Idle       Location
*130 vty 0     cisco      idle                 00:00:00   151.10.34.4
```

If the return traffic from R5 to R4 is allowed to be load-balanced between R1 and R3, Network Address Translation (NAT) will break because R1 has no configuration that is relevant to NAT.

To prevent this from happening, make a routing adjustment to RIP routes on R1 or R3. For example:

R1

```
router rip
 offset-list RIP-offset out 2 e0/1.20
!
ip access-list standard RIP-offset
 remark Manipulate NAT source networks, so R3 is preferred exit out of RIP
domain
 permit 151.10.43.0
 permit 151.10.34.0
```

On R4, issue the **telnet** command toward 151.10.10.1 port 80. Make sure that you establish the connection to R5.

```
R4#telnet 151.10.10.1 80
Trying 151.10.10.1, 80 ... Open

------------------------------------------------
Cisco 360 R&S Exercise Workbook
Product, POD location:  cierswbv5-ce-lab04-sc, SJ
Device:                 R5
------------------------------------------------

User Access Verification

Username: cisco
Password:

R5#who
    Line       User       Host(s)              Idle       Location
   0 con 0                 idle                 00:01:13   cierswbv5-ce-lab04-sc,
SJ
*  2 vty 0     cisco      idle                 00:00:00 151.10.34.4

  Interface    User               Mode         Idle     Peer Address

R5#exit

[Connection to 151.10.10.1 closed by foreign host]
R4#
```

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. Multicast

Issue: Enable multicast routing between R1 and R4, R4 and R3, and R3 and R5.

Solution:

Make sure that these Protocol Independent Multicast (PIM) configurations do not create any Reverse Path Forwarding (RPF) lookup problems.

Issue: Make R3 the root of the shared tree. Accomplish this task by configuring only R3.

Solution:

Configure R3 as both a Bootstrap Router Protocol (BSR) router candidate and a rendezvous point (RP) candidate.

This configuration is accomplished by entering the following two global configuration commands: **ip pim bsr-candidate Loopback103** and **ip pim rp-candidate loopback 103**.

On all multicast routers, make sure that RP is set up correctly:

```
R5#show ip pim rp mapping
PIM Group-to-RP Mappings

Group(s) 224.0.0.0/4
  RP 151.10.100.3 (?), v2
    Info source: 151.10.100.3 (?), via bootstrap, priority 0, holdtime 150
         Uptime: 01:52:56, expires: 00:01:59
R5#
```

Issue: Send multicast routing protocol hello messages five times a second between R1 and R4.

Solution:

The Multicast Sub-Second Convergence feature, one of many convergence improvements introduced into Cisco IOS Software, enables the following:

- Triggered RPF checks and improved periodic RPF checks
- Improved Internet Group Management Protocol (IGMP) and PIM state maintenance through new timer management techniques
- Improved scaling of the Multicast Source Discovery Protocol (MSDP) Source-Active (SA) cache

This task only requires modification on the PIM hello interval.

On R1 and R4, configure:

```
interface Ethernet0/2
 ip pim query-interval 200 msec
```

Issue: Ping the multicast group 229.20.20.20 from R1 and verify replies from each client router.

Solution:

Be on the lookout for RPF lookup problems when you perform this ping. By performing this ping, R1 is acting as a multicast source. Routers R4, R3, and R5 should respond.

Multicast Diagram

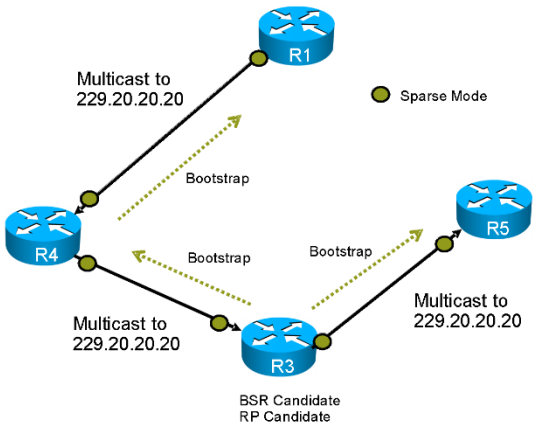

Verify a ping from R1:

```
R1#ping 229.20.20.20 source e0/2
Type escape sequence to abort.
Sending 1, 100-byte ICMP Echos to 229.20.20.20, timeout is 2 seconds:
Packet sent with a source address of 151.10.124.1

Reply to request 0 from 151.10.100.4, 1 ms
Reply to request 0 from 151.10.100.5, 1 ms
Reply to request 0 from 151.10.100.3, 1 ms
Reply to request 0 from 151.10.100.5, 1 ms
R1#
```

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**.