360RSW05-WKB-CA08

Cisco 360 CCIE R&S Exercise Workbook Introduction

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

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

Cisco 360 CCIE R&S Exercise Workbook Lab 8 Configuration Section

Table of Contents

Activity Objectives

When performing any Practice Lab, it is recommended that you formulate a test-taking strategy that includes the following activities. Some of these activities should be conducted in the actual lab:

- Download the latest copy of a Practice Lab, and then print it and read it carefully from beginning to end.
- Create a strategy for how to perform a Practice Lab.
- Draw diagrams if necessary.
- Create a checklist of general best practices to follow during the Practice Lab.
- Develop skill in finding issues in the lab so that you are able to uncover the hidden and complex internetworking issues.
- Carefully track your time so that you can develop good time-management techniques.
- Estimate the points that you have gained or lost to see where you are in your overall goal.

General Lab Instructions

Read the following instructions carefully. It is important to remember that if you misinterpret any directions, you could lose points. After you have read the "General Lab Instructions" section, read through the entire lab and look for connections between the tasks. Pay close attention to the "Restrictions and Goals" section because the information may reduce the configuration options that are available to you.

- Your pod should be cabled according to the example in the "Ethernet Switched Cabling Topology" figure and the IPv4 and IPv6 IGP diagrams.
- Each router should have an initial IP configuration loaded.
- You should be able to access all devices on your learner virtual pod via Telnet.
- To begin, check the following base configuration for each router and switch:
  - Configure a hostname on each device.
  - If a DNS server is being used in your pod, disable the DNS lookups.
  - Familiarize yourself with any Cisco IOS Software shortcuts.
  - Remember that some Cisco IOS command parameters and regular expressions are case-sensitive.
- Verify the following information on each router and switch:
  - Determine the Cisco IOS Software versions that are being used for the routers and the switches.
  - Verify that all the software on the routers and switches sees all physical interfaces.
- Review all the tasks in the scenario.

Difficulty Levels

Tasks are categorized as follows:

- **Basic:** These fundamental tasks are generally those tasks that are needed to provide the basic functions of the protocol or feature. You must complete these tasks to provide reachability and to move forward in the lab.
- **Intermediate:** These tasks include protocol features like routing optimization, route filtering, optimal path selection, load sharing, and summarization. Failure to complete these tasks will usually not affect later lab sections.
- **Advanced:** This category includes new Cisco IOS Software features and IP services, complex optimizations, and fine-tuning.

Scenarios are categorized as follows based on task classifications:

- Basic
- Basic to Intermediate
- Intermediate
- Intermediate to Advanced
- Advanced

Exercise Workbook Lab 8 Configuration Section

Grading and Duration

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

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

Difficulty Level

- Difficulty: Intermediate to Advanced

Restrictions and Goals

Note Read this section carefully.

- To receive credit for a subsection, you must fully complete the subsection as the requirements. You will *not* receive partial credit for partially completed subsections.
- IPv4 subnets that are displayed in the scenario diagram belong to network 160.20.0.0/16.
- IPv6 networks that are used in the scenario use a FEC0::/9 network.
- *Points will be deducted from multiple sections for failing to assign correct IPv4 addresses.*
- Do not use any static routes.
- Advertise loopback interfaces with their original masks for IPv4 and IPv6 protocols.
- Network 0.0.0.0/0 should not appear in any routing table (**show ip route**) , unless the instructions explicitly specify otherwise..
- Do not use the **ip default-network** command.
- All IP addresses involved in this scenario must be reachable, unless the instructions explicitly specify otherwise.
- Networks 192.50.*.* are excluded from the previous requirement to the extent that is explicitly described in this scenario.
- Unless the instructions explicitly specify otherwise, addresses and networks that advertised in the Border Gateway Protocol (BGP) section need to be reachable by all BGP routers but do not have to be reachable by routers that use only interior gateway protocol (IGP).
- Do not create new interfaces to fulfill IGP requirements, and do not create any summaries, unless the summary is required to meet explicitly stated scenario requirements.

- Do not introduce any new IPv4 or IPv6 addresses unless the instructions explicitly specify otherwise.
- Use only conventional routing algorithms.
- Do not modify the hostname, console, or vty configuration unless you are specifically asked to do so.
- Do not modify the initial interface or IP address numbering.

IPv4 IGP Diagram

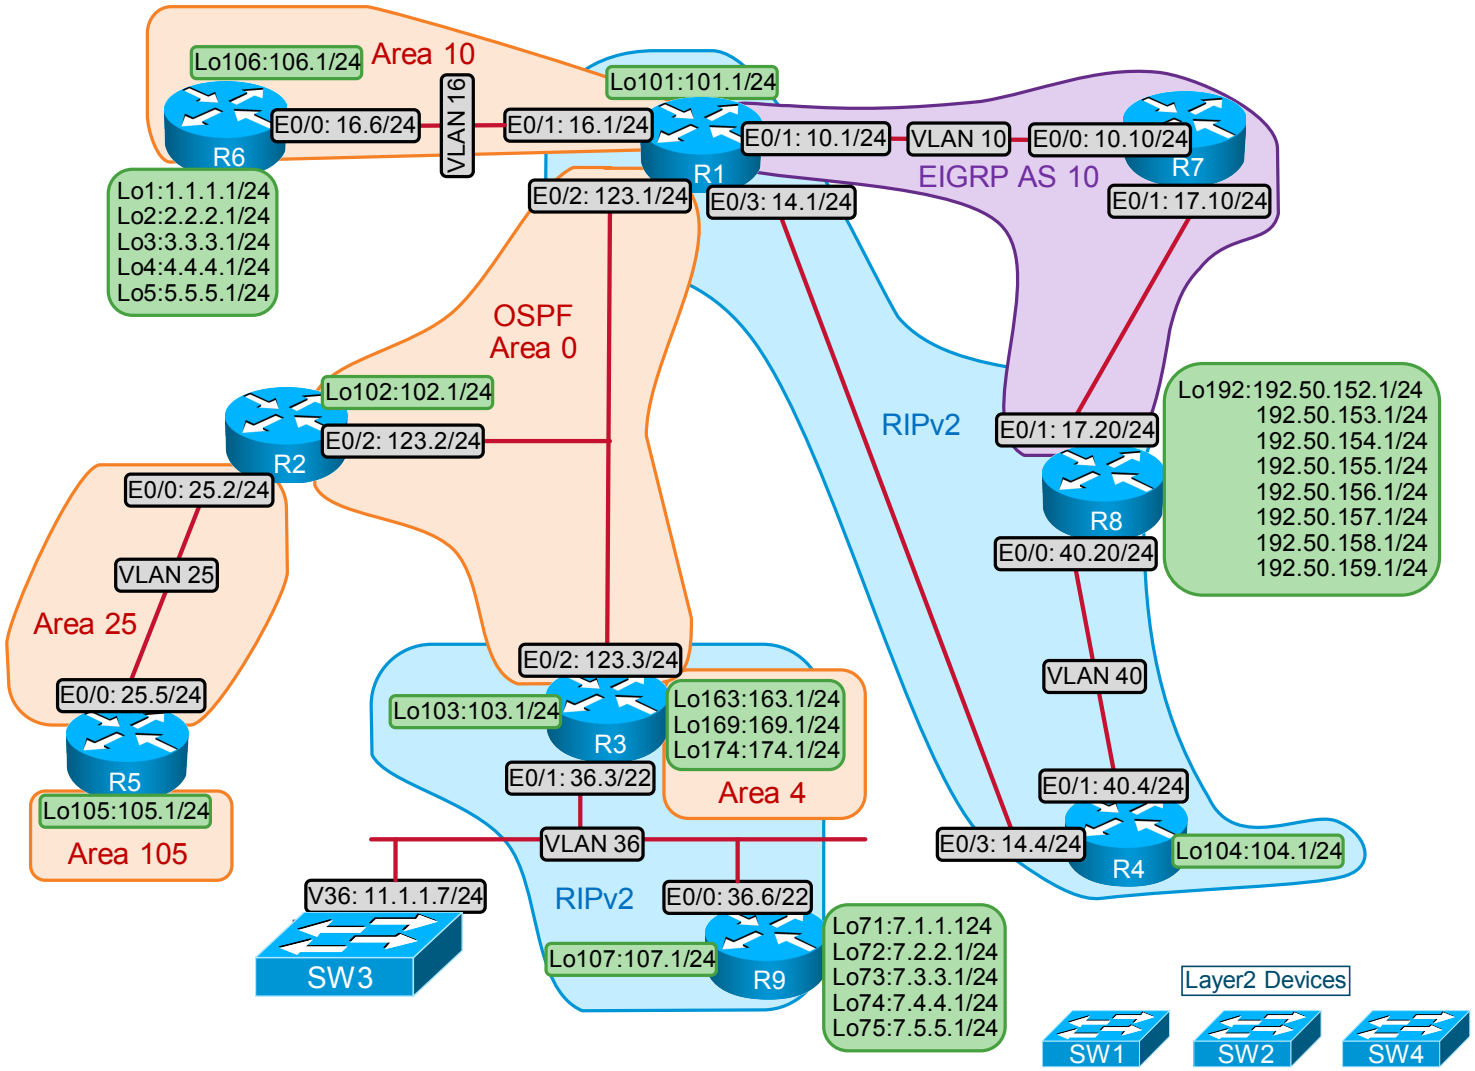

Ethernet Switched Cabling Topology

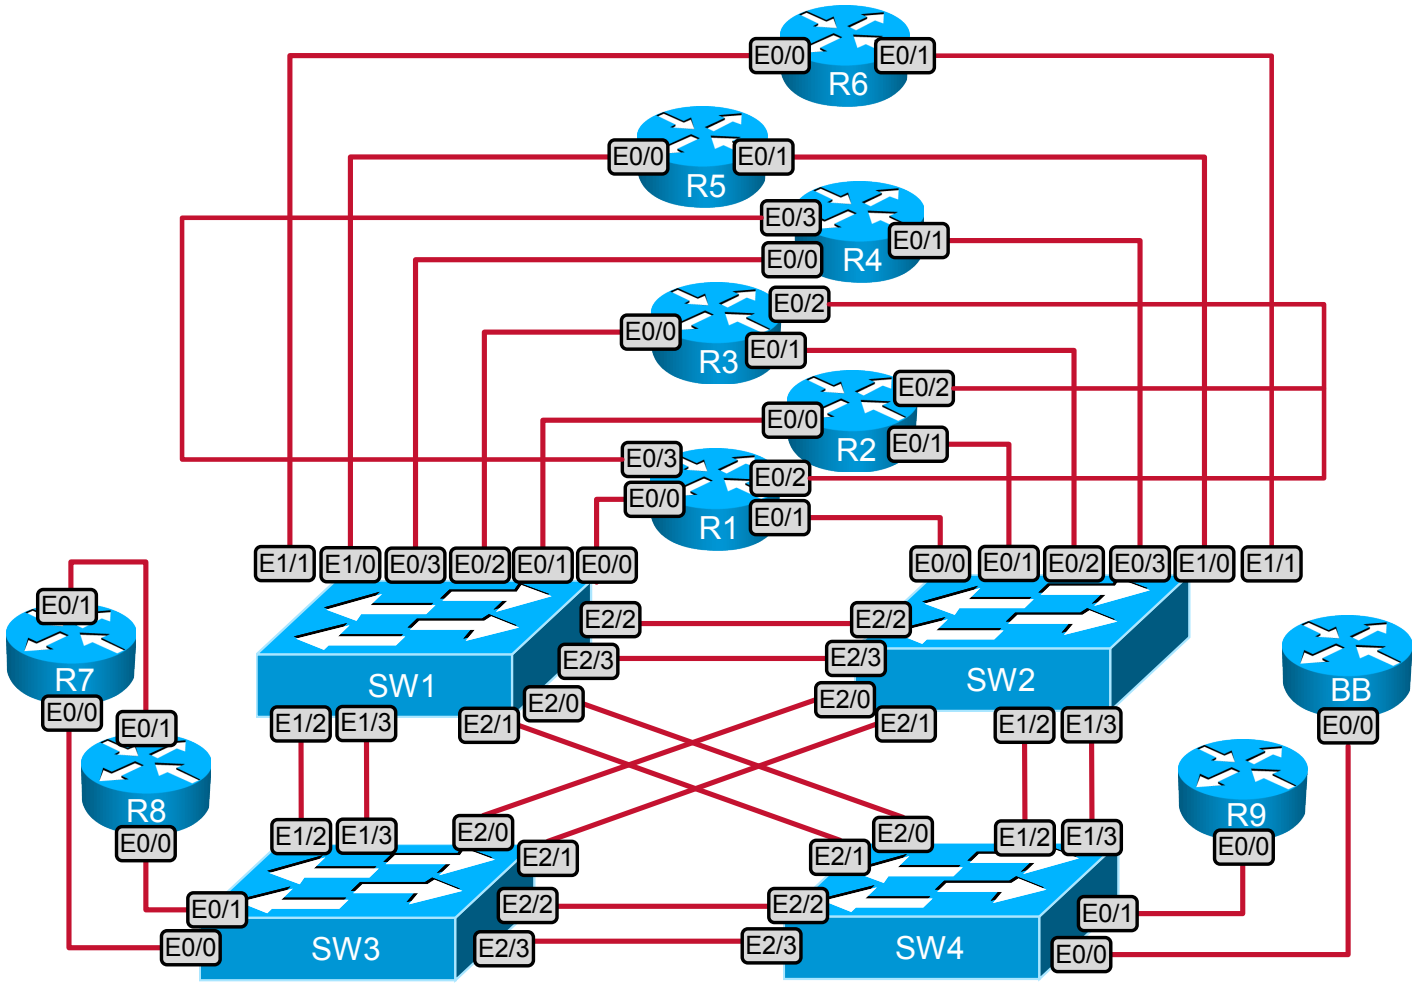

IPv6 IGP Diagram

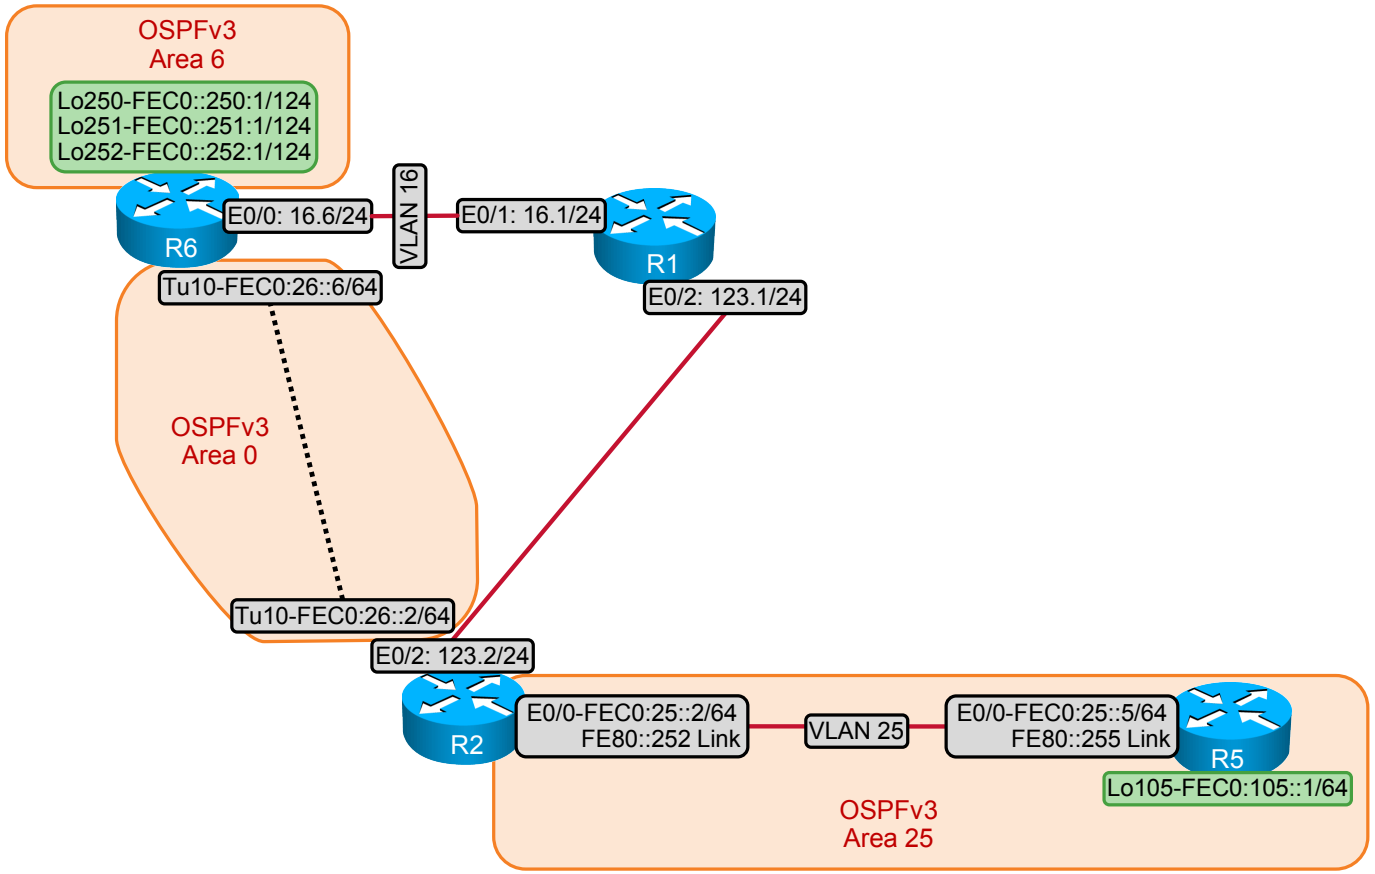

1. Switch Configuration Section (Total: 6 points)

1.1. Configure VLANs (Basic: 1 point)

- On SW1 and SW2, create VLANs 10, 16, 25, 36, and 40.
- On SW3 and SW4, create the VLANs as required by other sections of this scenario.

1.2. Configure Switch-to-Router Ports (Basic: 2 points)

- Configure the following switch-to-router connections.

Switch-to-Router Connections

Switch	Router	VLAN
SW2	R1 0/1	10, 16
SW1	R2 0/0	25
SW2	R3 0/1	36
SW2	R4 0/1	40
SW1	R5 0/0	25
SW1	R6 0/0	16
SW3	R7 0/0	10
SW3	R8 0/0	40
SW4	R9 0/0	36

- Configure switch ports as access VLAN ports whenever possible. Otherwise, use trunks.
- For switch-to-router trunking, use protocol 802.1Q. Limit VLANs on the switch-to-router trunks to those that are required by the scenario.
- Create the necessary switched virtual interfaces (SVIs) on the switches and assign the IP addresses that are specified in the diagram.
- Create the necessary Ethernet logical subinterfaces on the routers and assign the IP addresses that are specified in the diagram.

1.3. Configure VTP (Basic: 1 point)

- Use a transparent VLAN Trunking Protocol (VTP) mode.

1.4. Control Switch-to-Switch Links (Basic: 2 points)

- Ports listed in the following table must be administratively shut down. Verify that they are shut down and make sure that they remain in the shutdown state.

Switch Ports Shut Down

Switch	Port
SW1	1/3
	2/0
	2/1
	2/3
SW2	1/3
	2/0
	2/1
	2/3

Switch	Port
SW3	1/3
	2/0
	2/1
	2/2
	2/3
SW4	1/3
	2/0
	2/1
	2/2
	2/3

- Configure interfaces on active interswitch links according to the following table:

Switch-to-Switch Connections

Switch	Port	Switch	Port	Encapsulation
SW1	2/2	SW2	2/2	dot1q
SW1	1/2	SW3	1/2	dot1q
SW2	1/2	SW4	1/2	dot1q
SW4	0/0	Backbone	--	dot1q

- Only VLANs that need to carry traffic between the switches should be allowed on the trunk links.
- Do not enable IP routing on SW3.

2. IPv4 OSPF Section (Total: 8 points)

Note All OSPF routers must be configured with only one OSPF PID. *Points will be deducted from multiple sections for failing to assign only one OSPF PID on each specified router.* Use your IGP diagram to help guide configuration.

2.1. Create OSPF Areas (Basic: 2 points)

- Configure the 160.20.123.0/24 network between routers R1, R2, and R3 as OSPF Area 0.
- Configure OSPF Area 10 on subnet 160.20.16.0/24.
- Configure OSPF Area 25 on subnet 160.20.25.0/24.

2.2. Advertise Networks into OSPF (Basic: 2 points)

- Advertise loopback subnets 160.20.163.0/24, 160.20.169.0/24, and 160.20.174.0/24 in Area 4 from R3. Summarize the three entries with the most efficient mask.
- Advertise the loopback subnet 160.20.105.0/24 in Area 105.
- Advertise the loopback subnet 160.20.106.0/24 in Area 10.

2.3. Establish OSPF Adjacencies (Intermediate: 1 point)

- Use the OSPF non-broadcast network type on the 160.20.123.0/24 network.
- R1 should be the DR. R2 and R3 should be the DROTHERs.

2.4. Control OSPF Routing (Intermediate: 2 points)

- Make sure that the routers in OSPF Area 10 possess the minimum amount of routing information to reach all destinations within your pod. The solution may include network 0.0.0.0/0 in the routing table of R6.

2.5. Verify Connectivity (Basic: 1 point)

- Verify that all OSPF prefixes that are specified in this section can be reached from all devices in the OSPF domain.

3. IPv4 RIP Section (Total: 6 points)

3.1. Enable RIP (Basic: 1 point)

- Configure RIPv2 over the link between routers R1 and R4.
- Configure RIPv2 between routers R4 and R8.
- Configure RIPv2 between routers R3 and R9.

3.2. Control RIP updates (Intermediate: 2 points)

- Make sure that RIP advertises only over the links in the previous task.

3.3. Advertise Networks into RIP (Basic: 1 point)

- On R8, advertise the Loopback192 interface networks into RIP without using a network statement:
  - 192.50.152.0/24
  - 192.50.153.0/24
  - 192.50.154.0/24
  - 192.50.155.0/24
  - 192.50.156.0/24
  - 192.50.157.0/24
  - 192.50.158.0/24
  - 192.50.159.0/24

3.4. Control RIP Routing (Intermediate: 2 points)

- Make sure that R8 advertises only the following networks to R4:
  - 192.50.153.0/24
  - 192.50.155.0/24
  - 192.50.157.0/24
  - 192.50.159.0/24

- Use the minimal number of ACL statements to match these networks.

4. IPv4 EIGRP Section (Total: 7 points)

4.1. Enable EIGRP (Basic: 1 point)

- Configure EIGRP AS 10 on subnet 160.20.10.0/24 between R1 and R7.
- Configure EIGRP AS 10 on subnet 160.20.17.0/24 between R7 and R8.

4.2. Advertise Networks into EIGRP (Basic: 2 points)

- On R8, advertise the following networks into EIGRP AS 10 without using a network statement:
  - 192.50.152.0/24
  - 192.50.153.0/24
  - 192.50.154.0/24
  - 192.50.155.0/24
  - 192.50.156.0/24
  - 192.50.157.0/24
  - 192.50.158.0/24
  - 192.50.159.0/24

4.3. Control EIGRP Routing Updates (Intermediate: 3 points)

- Make sure that R8 advertises only the following networks to R7:
  - 192.50.153.0/24
  - 192.50.155.0/24
  - 192.50.157.0/24
  - 192.50.159.0/24
- Use the minimal number of ACL statements to match these networks.

4.4. Verify Connectivity (Intermediate: 1 point)

- Verify that all EIGRP prefixes that are specified in this section can be reached from all devices in the EIGRP domain. Note that subnets 192.50.152.0/24, 192.50.154.0/24, 192.50.156/24, and 192.50.158.0/24 are excluded from all reachability requirements.

5. IPv4 Route Redistribution Section (Total: 4 points)

5.1. Obtain Universal Connectivity (Advanced: 2 points)

- On R1, mutually redistribute EIGRP AS 10 and OSPF.
- Also on R1, redistribute RIP into EIGRP and OSPF.
- On R3, mutually redistribute RIP and OSPF.
- Do not perform any other redistribution, except as specified on R8 in the RIP and EIGRP sections.

5.2. Complete Redistribution Tuning (Intermediate: 2 points)

- Fulfill the reachability requirements without allowing 0.0.0.0/0 in any routing table, other than R6.

6. BGP Section (Total: 8 points)

6.1. Configure Processes and Peers (Intermediate: 2 points)

- Configure BGP AS 100 on R1, R2, R3, and R4.
- Configure BGP AS 600 on R6.
- Configure BGP AS 700 on R9.
- Configure BGP peer relationships between AS 100 and AS 600 using peers R1 and R6.
- Configure BGP peer relationships between AS 100 and AS 700 using peers R3 and R9.
- Do not allow a full mesh of Internal Border Gateway Protocol (IBGP) peer relationships within AS 100. Provide redundant Network Layer Reachability Information (NLRI) exchange using routers R2 and R3.

6.2. Advertise BGP Prefixes (Intermediate: 2 points)

- Advertise the following networks in the AS 600 from R6:
  - 1.1.1.0/24
  - 2.2.2.0/24
  - 3.3.3.0/24
  - 4.4.4.0/24
  - 5.5.5.0/24
- Advertise the following networks in AS 700 from R9:
  - 7.1.1.0/24
  - 7.2.2.0/24
  - 7.3.3.0/24
  - 7.4.4.0/24
  - 7.5.5.0/24

6.3. Control BGP Routing (Intermediate: 4 points)

- Configure R1 and R3 so that AS 100 accepts from its External Border Gateway Protocol (EBGP) peers only those prefixes with a third octet of 4 or 5 that originated from the connected AS (AS path length of one AS number). Filtered networks are excluded from the reachability requirement.
- Your filtering configuration should be the same on R1 and R3.

7. Router Maintenance Section (Total: 4 points)

7.1. Complete Address Administration (Intermediate: 4 points)

- Configure the 11.1.1.3/24 address on R3's Ethernet0/1 interface without changing any pre-existing IP addresses.
- Configure the 11.1.1.7/24 address on SW3's VLAN 36 interface.

- Do not advertise subnet 11.1.1.0/24 into any routing protocol.
- Ensure that workstations on the 11.1.1.0/24 private address space connected to VLAN 36 can reach the rest of the network using a portion of the address space of the R3 160.20.36.0/22 subnet.
- SW3 should be reachable from the rest of the network using the 160.20.36.130 address.

8. IPv6 Routing Section (Total: 10 points)

8.1. Configure IPv6 Interfaces and Link-Local Addresses (Basic: 2 points)

- Configure IPv6 link-local addresses for interfaces on subnet 160.20.25.0/24 according to the following table:

IPv6 Link-Local Address Assignment

Router	Link	IPv6 Link-Local Address
R2	160.20.25.2	FE80::252
R5	160.20.25.5	FE80::255

- Verify connectivity using configured IPv6 link-local addresses.

8.2. Configure IPv6 Addresses (Intermediate: 2 points)

- Configure a logical point-to-point link between R2 and R6 to carry IPv6. Do not enable IPv6 on R1.
- Configure IPv6 addresses according to the following table:

IPv6 Address Assignment

Router	Link	IPv6 Address	OSPF Area
R2	160.20.25.2	FEC0:25::2/64	Area 25
	Tunnel with source 160.20.123.2	FEC0:26::2/64	Area 0
R5	160.20.105.1	FEC0:105::1/64	Area 25
	160.20.25.5	FEC0:25::5/64	Area 25
R6	Lo250 (no IPv4)	FEC0::250:1/124	Area 6
	Lo251 (no IPv4)	FEC0::251:1/124	Area 6
	Lo252 (no IPv4)	FEC0::252:1/124	Area 6
	Tunnel with source 160.20.16.6	FEC0:26::6/64	Area 0

8.3. Configure OSPFv3 (Basic: 2 points)

- Use the OSPFv3 nonbroadcast multiaccess (NBMA) network on subnet FEC0:25::/64.
- Make sure that routers R2 and R5 see only a summary route representing the IPv6 loopback interfaces on R6.

8.4. Configure IPv6 QoS (Intermediate: 2 points)

- All traffic that leaves R6 destined to IPv6 addresses should be tagged with the differentiated services code point (DSCP) value Expedited Forwarding (EF). Do not change the IPv6 header. The configuration should be applied on R6 and should not use the Modular QoS CLI (MQC), committed access rate (CAR), or policy-based routing (PBR).

8.5. Configure QoS Monitoring (Intermediate: 2 points)

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

9. QoS Section (Total: 5 points)

9.1. Configure Traffic Management (Intermediate: 5 points)

- Assume there is a traffic generator on VLAN 10 connected via SW4 E0/0. It is generating five UDP packets per second. The packet size is 1024 bytes. The UDP stream is destined to 160.20.10.1 port 5111.
- Limit incoming UDP traffic that is destined to port 5111 to a rate of 8000 b/s on the router R1 interface on subnet 160.20.10.0/24.
- Configure the minimal values for burst size and extended burst size. The solution should continue to provide effective limiting if the packet size on the traffic generator is changed.
- Drop excessive traffic.

10. Network Security Section (Total: 5 points)

10.1. Configure Secure Copy Protocol (Intermediate: 5 points)

- Configure R2 as a Secure Copy Protocol (SCP) server in testlab.com. Only the locally configured user noc with the password cisco is permitted to use this feature.
- Configure vty lines 0 through 4 to permit only SSH traffic, and only traffic from the IP address 160.20.25.5. The remaining vty lines should not permit any connections.
- Test the configuration by reading R2's running configuration from R5.

11. Switch Specialties Section (Total: 6 points)

11.1. Configure Interface (Intermediate: 3 points)

- Set port E1/1 of SW2 to bypass the learning and listening states of spanning tree.

11.2. Complete MAC Address Administration (Intermediate: 3 points)

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

12. Multicast Section (Total: 7 points)

12.1. Configure PIM (Intermediate: 1 point)

- Enable multicast routing between routers R1, R2, R3, R5, and R6.
- Enable a multicast routing protocol that will use any unicast routing protocol for source address determination and that is also based on a shared tree.

- Make R1 the root of the shared tree. Do not use any dynamic methods to discover or advertise the root of the shared tree.

12.2. Configure IGMP (Intermediate: 2 points)

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

12.3. Complete IGMP Tuning (Intermediate: 2 points)

- Configure R3 so that a workstation on VLAN 36 could join only multicast group 227.7.7.7 and not any other groups.

12.4. Verify Multicast Connectivity (Advanced: 2 points)

- Ping the multicast group 227.7.7.7 from R4 using the source address 160.20.14.4. Verify that responses are received from clients on R1, R2, R3, R5, and R6.