360RSW05-WKB-CA04

Cisco 360 CCIE R&S Exercise Workbook Introduction

The Cisco 360 CCIE® R&S Version 5 Exercise Workbook contains 20 challenging scenarios at the Cisco CCIE level that can be used for rigorous self-paced practice. The Exercise Workbook scenarios include both a troubleshooting section and a configuration section.

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

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 carefully 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" diagram and the IPv4 and IPv6 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 virtual switches.
- 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 4 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 per the requirements. You will *not* receive partial credit for partially completed subsections.
- Partial 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**, or **ip default-network** commands.
- Do not introduce any new IP addresses unless specifically directed to do so.
- 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 or 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.

IPv4 IGP Diagram

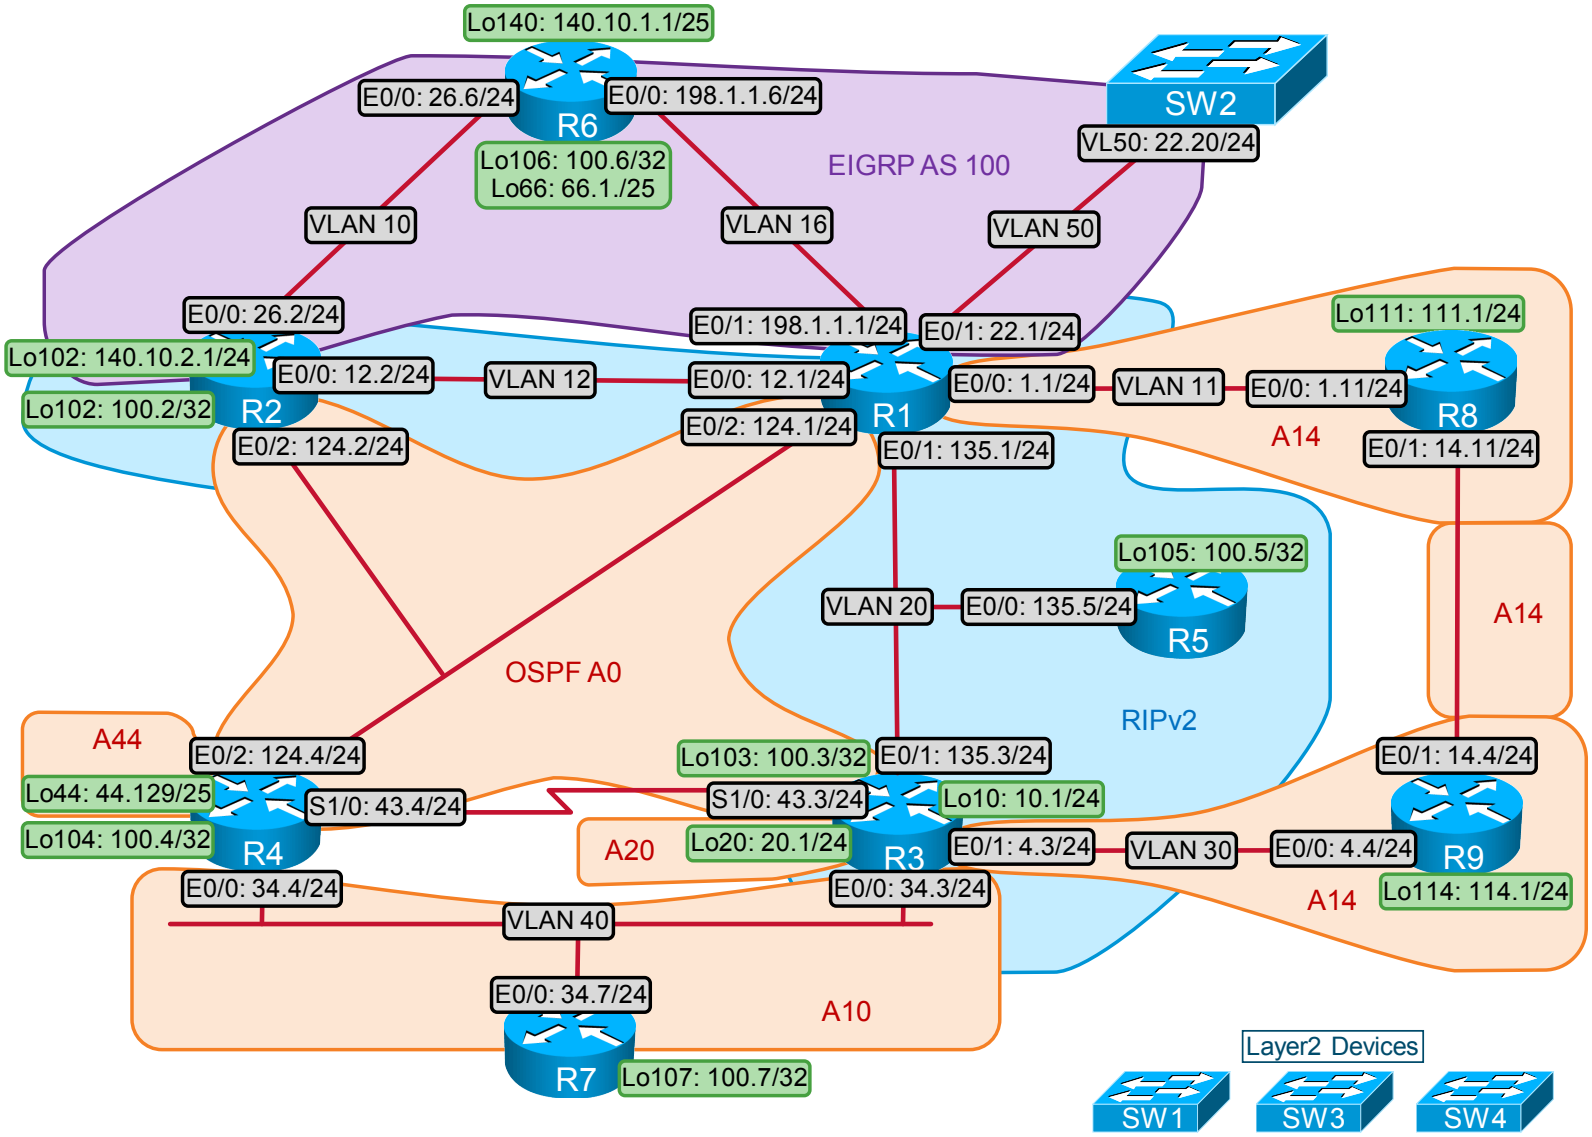

Ethernet Switched Cabling Topology

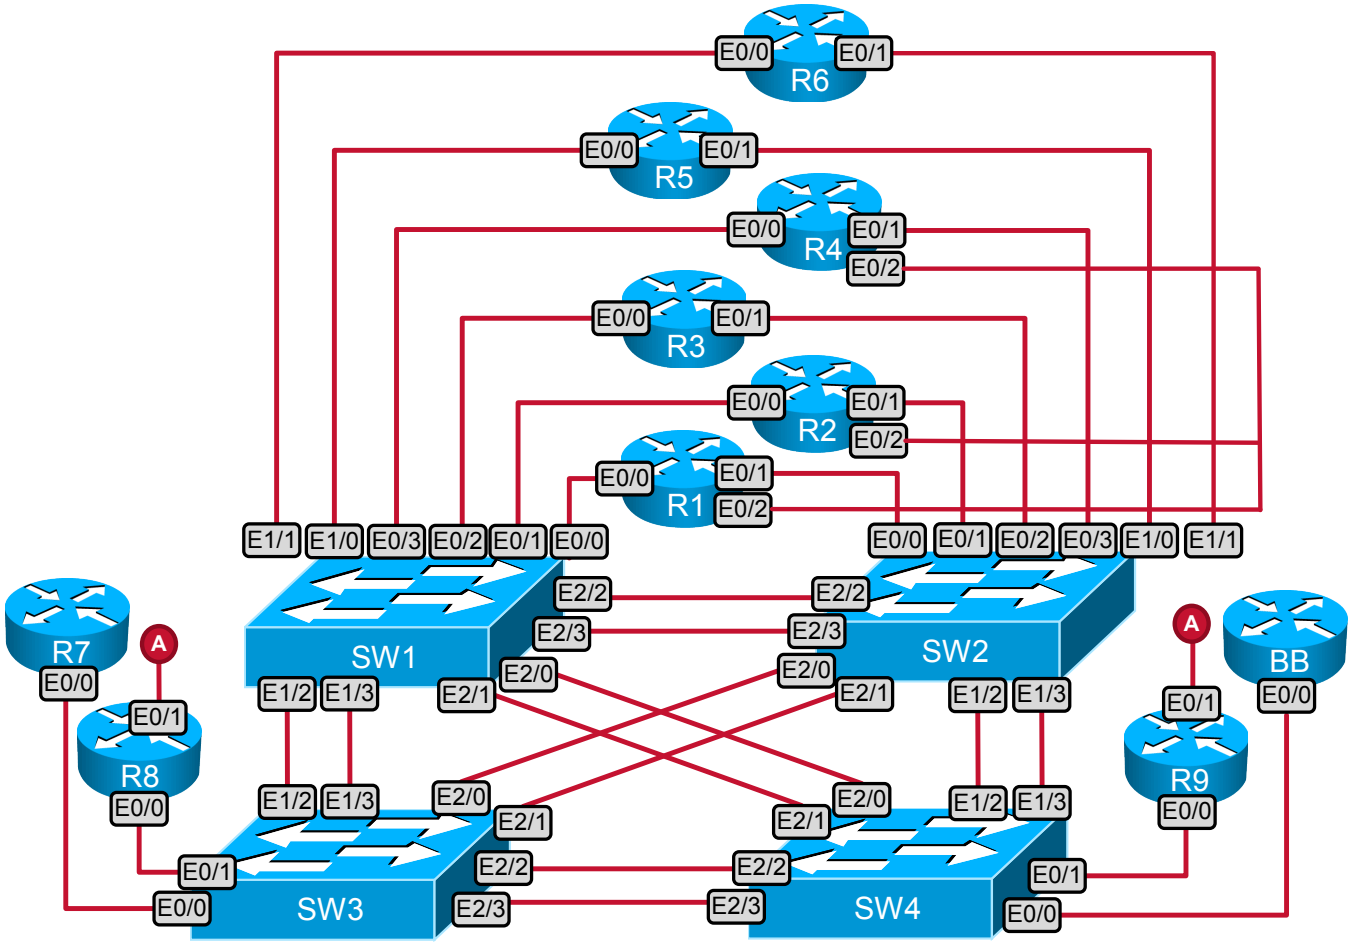

IPv6 Topology Diagram

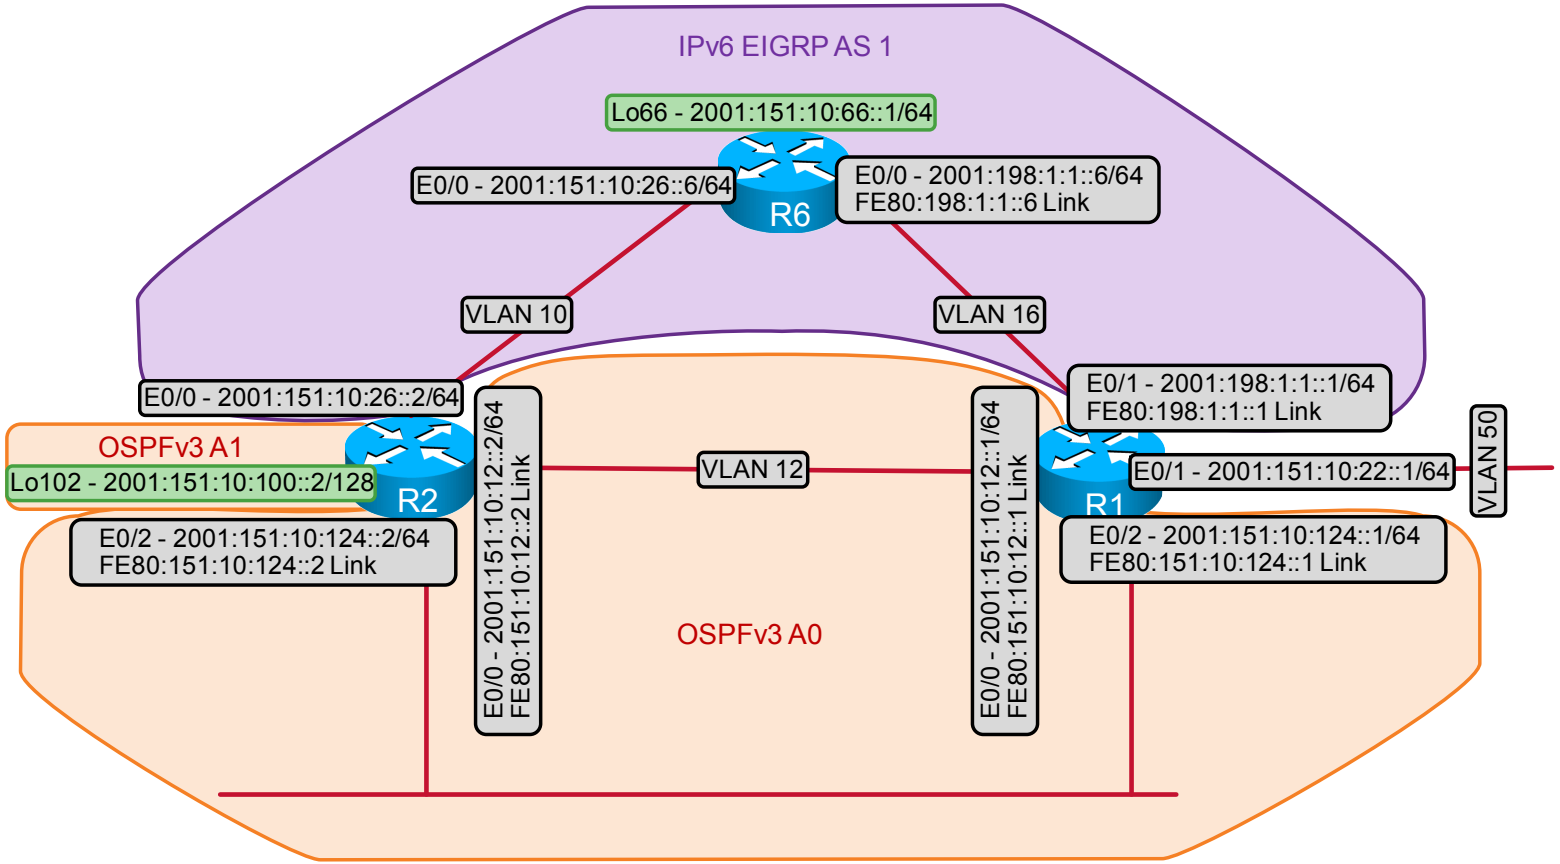

1. Switch Configuration Section (Total: 8 points)

Note Port 0/0 on SW4 is connected to the backbone. The configuration of this port should be left as the default. Do not change it or the connectivity with the backbone will be lost. Healthy trunk status is displayed as shown here:

```
Mode         Encapsulation  Status
on           802.1q         trunking
```

Do not change any initially configured link speeds.

1.1. Configure VLANs (Basic: 2 points)

- Create the VLANs referenced in the following table:

VLANs

VLAN	VLAN NAME
VLAN 10	Net-10
VLAN 11	Net-11
VLAN 12	Net-12
VLAN 16	Net-16
VLAN 20	Net-20
VLAN 30	Net-30
VLAN 40	Net-40
VLAN 50	Net-50

- Configure only the necessary VLANs on each switch (read other sections).
- Configure the following switch-to-router connections:

Switch-to-Router Connections

Switch	Router	VLAN	Mode
SW1	R1	VLANs 11 and 12	Trunk dot1q
SW1	R2	VLANs 10 and 12	Trunk dot1q
SW1	R3	VLAN 40	Access
SW1	R4	VLAN 40	Access
SW1	R5	VLAN 20	Access
SW1	R6	VLANs 10 and 16	Trunk dot1q
SW2	R1	VLANs 16, 20, and 50	Trunk dot1q
SW2	R3	VLANs 20 and 30	Trunk dot1q
SW3	R7	VLAN 40	Access
SW3	R8	VLAN 11	Access
SW4	R9	VLAN 30	Access

- Create the necessary switch virtual interfaces (SVIs) and assign the IP addresses that are specified in the diagram.

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

- Configure interfaces on active switch-to-switch links according to the following table, and verify that ports that indicate they are administratively shut down remain in the shutdown state:

Switch-to-Switch Connections

Switch/Port		Switch/Port		Mode
SW1	1/2	SW3	1/2	Administratively shut down
	1/3		1/3	Trunk dot1q
	2/0	SW4	2/0	Trunk dot1q
	2/1		2/1	Administratively shut down
	2/2	SW2	2/2	Trunk dot1q
	2/3		2/3	Trunk dot1q
SW2	1/2	SW4	1/2	Administratively shut down
	1/3		1/3	Access
	2/0	SW3	2/0	Administratively shut down
	2/1		2/1	Administratively shut down
SW3	2/2	SW4	2/2	Trunk dot1q
	2/3		2/3	Trunk dot1q

- Do not allow switches to advertise VLANs.

1.3. Tune Switch-to-Switch Links (Intermediate: 2 points)

- Allow only the necessary VLANs on the trunks between SW1 and SW2, between SW1 and SW3, between SW1 and SW4, and between SW3 and SW4.

1.4. Tune STP (Intermediate: 2 points)

- 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.
- When sending frames for VLAN 40, make sure that SW1 prefers interface 2/0 to interface 1/3.
- SW4 should load-balance between interfaces 2/2 and 2/3.

2. IPv4 OSPF Section (Total: 6 points)

2.1. Create OSPF Areas (Basic: 2 points)

- Place all Ethernet0/2 interfaces in Area 0.
- Choose the OSPF network type that will form a collection of point-to-point adjacencies on subnet 151.10.124.0/24 and will generate host entries for the Ethernet0/2 interfaces.
- Choose the OSPF nonbroadcast network type on the serial link between R3 and R4.

- Configure OSPF Area 10 on the subnet between R3, R4, and R7.
- Assign loopback 151.10.20.1/24 to Area 20 on R3.
- Assign loopback 151.10.44.129/25 to Area 44 on R4.

2.2. Control OSPF Advertisements (Intermediate: 2 points)

- 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.
- Introduce loopback 151.10.100.4/32 on R4 and loopback 151.10.100.7/32 on R7 into OSPF without using a network statement or interface configuration OSPF statements.

2.3. Tune OSPF (Intermediate: 2 points)

- On Area 10, make R7 the preferred candidate for the designated router (DR) and R4 the preferred candidate for backup designated router (BDR).
- Make sure that R4 can reach network 151.10.20.0 using the OSPF path across VLAN 40.
- Verify that all OSPF prefixes that are specified in this section can be reached from all devices in the OSPF domain.

3. IPv4 EIGRP Section (Total: 4 points)

3.1. Configure EIGRP AS 100 (Basic: 2 points)

- Configure EIGRP AS 100 between R1 and R6, R2 and R6, and R1 and SW2.
- Do not use any summarization in this domain.
- Advertise the following loopback interfaces on R6 as internal Enhanced Interior Gateway Routing Protocol (EIGRP) networks:
  - 151.10.66.1/25
  - 151.10.100.6/32
- Advertise the following loopback interface on R2 as an internal EIGRP network:
  - 140.10.2.1/24
- Advertise the following loopback interface on R6. This route should show as “D EX” in the R2 routing table:
  - 140.10.1.1/25

3.2. Tune EIGRP (Intermediate: 2 points)

- Prevent EIGRP queries from being sent to SW2. SW2 should be able to ping the rest of the network.
- 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.
- Verify that all EIGRP prefixes that are specified in this section can be reached from all devices in the EIGRP domain.

4. IPv4 RIP Section (Total: 2 points)

4.1. Enable RIP and Control Updates (Basic: 2 points)

- Configure Routing Information Protocol version 2 (RIPv2) only between R1, R2, R3, and R5.
- Make RIP send updates three times more often than they are sent by default.

5. Redistribution Section (Total: 5 points)

5.1. Obtain Universal Connectivity (Advanced: 2 points)

- Perform a mutual redistribution of dynamic IGPs between the following:
  - RIP and OSPF on R1 and R3
  - RIP and EIGRP on R1
  - OSPF and EIGRP on R1
- Redistribute RIP into EIGRP and EIGRP into OSPF on R2.
- Perform the **redistribute connected** command where required and not restricted by the scenario.
- Subnets 140.10.1.0/25 and 140.10.2.0/24 should be advertised into OSPF on R2.

5.2. Verify Connectivity (Advanced: 3 points)

- Verify that all IPv4 IGP prefixes that are specified on the IPv4 IGP diagram can be reached from all devices. See the "Restrictions and Goals" section.

6. BGP Section (Total: 7 points)

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

- Use the synchronization method in this section.
- Configure BGP autonomous systems (AS) according to the following table:

BGP AS Assignment

Device	AS
R2	20
R1	134
R3	134
R5	50
SW2	52

- Configure peering between these devices:
  - R1 and R2
  - R3 and R5
  - R1 and SW2

6.2. Advertise BGP Prefixes (Intermediate: 2 points)

- Add network 140.10.1.0/25 to the R2 BGP routing process.

- Advertise the network 140.10.2.0/24 on R2.
- Make sure that networks 140.10.1.0/25 and 140.10.2.0/24 reside in the BGP table on R5.
- Both networks 140.10.1.0/25 and 140.10.2.0/24 should be added to BGP on R2 without using redistribution commands.

6.3. Set Community Tags (Advanced: 2 points)

- 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 should see the prefixes marked with communities.

6.4. Tune BGP Adjacency (Advanced: 1 point)

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

7. MPLS Layer 3 VPN Section (Total: 13 points)

7.1. Configure VRFs on R1 and R3 (Basic: 3 points)

- Configure a VRF called VPN 14 on R1 and R3. Use route distinguisher value 134:14.
- Assign interface E0/0.11 on R1 and interface E0/1.30 on R3 to this VRF.

7.2. Use OSPF for Provider-Edge-to-Customer-Edge Routing (Basic: 3 points)

- Configure OSPF Area 14 on all VPN14 interfaces, including all of the interfaces on R8 and R9.

7.3. Configure BGP to Carry VPN Routes between R1 and R3 (Basic: 3 points)

- Enable the exchange of VPNv4 addresses between R1 and R3 and the transmission of VPN 14 traffic across VLAN 20.
- Mutually redistribute BGP and OSPF on R1 and R3 so that VPN routes will be advertised from site to site.

7.4. Optimize Site-to-Site Traffic (Intermediate: 4 points)

- The Ethernet connection between R8 and R9 is intended as a backup route. It should be configured with an OSPF cost of 100. Without changing any OSPF cost values, make sure that the primary path between R8 and R9 is across the MPLS core. You are permitted to add two interfaces and two IP addresses to the VRF to meet this requirement.

Note	IPv4 addresses in VPN 14 do not have to be reachable from outside the VPN.

8. IPv6 Routing Section (Total: 9 points)

8.1. Configure IPv6 EIGRP AS 1 and IPv6 OSPF (Intermediate: 3 points)

- IPv6 addresses should be configured according to the table:

IPv6 Address Assignment

Router	IPv4 Interface	Link-Local Addresses	Routable Addresses
R1	VLAN 50	Default	2001:151:10:22::1/64
R1	E0/2	FE80:151:10:124::1	2001:151:10:124::1/64
R1	VLAN 12	FE80:151:10:12::1	2001:151:10:12::1/64
R1	VLAN 16	FE80:198:1:1::1	2001:198:1:1::1/64
R2	Loopback 102	Default	2001:151:10:100::2/128
R2	E0/0	Default	2001:151:10:26::2/64
R2	E0/2	FE80:151:10:124::2	2001:151:10:124::2/64
R2	VLAN 12	FE80:151:10:12::2	2001:151:10:12::2/64
R6	Loopback 66	Default	2001:151:10:66::1/64
R6	E0/0	Default	2001:151:10:26::6/64
R6	VLAN 16	FE80:198:1:1::6	2001:198:1:1::6/64

- Verify connectivity between IPv6 addresses on the Ethernet links.
- Configure IPv6 EIGRP AS 1 between R1 and R6, and between R2 and R6.
- 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.
- Assign the R2 loopback 102 interface to OSPF Area 1. All other OSPF routers should see this prefix summarized to /64.

8.2. Tune IPv6 OSPF (Advanced: 2 points)

- 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 to be missed before the neighbor is declared inactive.
- Assuming that the OSPFv3 network topology for IPv6 is stable, suppress unnecessary flooding of link state advertisements (LSAs) on the connection between the Ethernet0/2 interface of R1 and R2.

8.3. Configure IPv6 Services (Advanced: 2 points)

- Configure R1 to use an IPv6 DNS server that is located at address 2001::1. Do not enable DNS-based name-to-address translation.
- Allow R1 to use Telnet to connect to loopback 102 of R2 using the following:
  - The **R2-v6** name for the IPv6 protocol
  - The **R2** name for the IPv4 protocol
- 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.

8.4. Ensure Entire IPv6 Connectivity (Intermediate: 2 points)

- 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 redundantly, as follows:
  - The primary path should be via the link between the Ethernet0/2 interfaces.

- The secondary path should be via VLAN 12.
- The tertiary path should be via R6.

- Make sure that all global IPv6 addresses that are specified in the "IPv6 Routing" section are reachable. Perform IPv6 redistribution as necessary.

9. QoS Section (Total: 6 points)

9.1. Mark Traffic (Intermediate: 2 points)

- Traffic that is sent by SW2 to BGP networks in AS 20 should be marked on R1 with the following:
  - IP precedence 3 for prefixes that are marked with BGP community 20:3
  - IP precedence 5 for prefixes that are marked with BGP community 20:5

9.2. Control Traffic on the Ethernet0/2 interface on R1 (Advanced: 4 points)

- Configure Quality of Service (QoS) on R1.
- The R1 Ethernet0/2 interface leads to another QoS domain. Modify QoS marking on the packets that are leaving the interface as follows:

QoS Precedence-to-DSCP Modification

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

- 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.

10. System Administration Section (Total: 8 points)

10.1. Monitor CPU Utilization (Advanced: 3 points)

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

10.2. Secure Access (Advanced: 3 points)

- 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 with name **R5.cisco.com** and minimal allowable modulus.
- Configure R2 to establish SSHv2 sessions to the loopback 0 interface on R5 when **connect-to-r5** is entered.

10.3. Configure Advanced Cisco IOS Software Features (Advanced: 2 points)

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

11. Address Administration Section (Total: 3 points)

11.1. Address Administration (Advanced: 3 points)

- 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.

12. Multicast Section (Total: 5 points)

12.1. Configure PIM (Basic: 2 points)

- Enable multicast routing between R1 and R4, R4 and R3, and R3 and R5. Enable a multicast routing protocol that will use PIMv2 messaging only to build the shared tree.
- Configure R3, R4, and R5 to join the multicast group 229.20.20.20. Associate this multicast group with a loopback interface on each router.
- Make R3 the root of the shared tree. Accomplish this task by configuring only R3.

12.2. Tune PIM Configuration (Intermediate: 2 points)

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

12.3. Verify Multicast Connectivity (Intermediate: 1 point)

- Ping the multicast group 229.20.20.20 from R1 IP address 151.10.124.1.

