IEWB-RS Volume 2 Lab 3

Difficulty Rating (10 highest): 6

Lab Overview:

The following scenario is a practice lab exam designed to test your skills at configuring Cisco networking devices. Specifically, this scenario is designed to assist you in your preparation for Cisco Systems' CCIE Routing & Switching Lab exam. However, remember that in addition to being designed as a simulation of the actual CCIE lab exam, this practice lab should be used as a learning tool. Instead of rushing through the lab in order to complete all the configuration steps, take the time to research the networking technology in question and gain a deeper understanding of the principles behind its operation.

Lab Instructions:

Prior to starting, ensure that the initial configuration scripts for this lab have been applied. For a current copy of these scripts, see the Internetwork Expert members' site at http://members.INE.com. If you have any questions related to the scenario solutions, visit our CCIE support forum at http://IEOC.com.

Refer to the attached diagrams for interface and protocol assignments. Any reference to X in an IP address refers to your rack number, while any reference to Y in an IP address refers to your router number.

Upon completion, all devices should have full IP reachability to all networks in the routing domain, including any networks generated by the backbone routers unless explicitly specified.

Lab Do's and Don'ts:

- Do not change or add any IP addresses from the initial configuration unless otherwise specified or required for troubleshooting
- If additional IP addresses are needed but not specifically permitted by the task use IP unnumbered
- Do not change any interface encapsulations unless otherwise specified
- Do not change the console, AUX, and VTY passwords or access methods unless otherwise specified
- Do not use any static routes, default routes, default networks, or policy routing unless otherwise specified
- Save your configurations often

Grading:

This practice lab consists of various sections totaling 79 points. A score of 64 points is required to pass the exam. A section must work 100% with the requirements given in order to be awarded the points for that section. No partial credit is awarded. If a section has multiple possible solutions, choose the solution that best meets the requirements.

Point Values:

The point values for each section are as follows:

Section	Point Value
Layer 2 Technologies	9
IPv4	28
IPv6	4
MPLS VPN	9
Multicast	6
Security	6
Network Services	10
QoS	7

GOOD LUCK!

1. Layer 2 Technologies

1.1 IP Bridging

- R1 and R3 are in the same IP subnet, but in different broadcast domains.
- Configure R6 to bridge IP traffic between VLAN 16 and VLAN 36.
- Ensure that the rest of the routing domain can communicate with both R1 and R3 via IP.

2 Points

1.2 Spanning-Tree Protocol

- Configure SW1 as the spanning-tree root for VLANs: 4, 44, 52, and 63.
- All traffic between SW1 and SW2 for these VLANs should transit the trunk between SW1 and SW2's port Fa0/15.
- In the case that port Fa0/15 goes down, traffic for these VLANs should transit port Fa0/14.
- As a last resort, traffic for these VLANs should transit port Fa0/13 if both of the other trunk links are down.
- This configuration should be done on SW1.

2 Points

1.3 Spanning-Tree Protocol

- In order to minimize network downtime in the event of a failure, configure SW2 so that traffic continues forwarding within three seconds if either port Fa0/15 or Fa0/14 goes down.
- This should be accomplished while running PVST.

3 Points

1.4 Switch Management

- Configure SW1 and SW2 to be managed via SNMP using the following parameters:
    - Contact: CCIE Lab SW1.
    - Location: San Jose, CA US.
    - Chassis ID: 221-787878.
- The network management station's IP address is 136.X.2.100, and it will be expecting the RO community string to be CISCORO and the RW community string to be CISCORW.
- SW1 and SW2 should generate SNMP traps for changes related to VTP, using the community string CISCOTRAP.

2 Points

2. IPv4

2.1 OSPF

- Ensure that R2 uses R5 as the next hop to reach R4, and vice versa.
- Advertise the Loopback 0 interfaces of R1, R2, R4, and R5 into OSPF area 0.
- These routes should appear with a subnet mask of /24 with the exception to the prefixes that you need to appear as /32.

2 Points

2.2 OSPF

- Configure OSPF area 45 on the PPP link between R4 and R5.
- This link will be used primarily as a backup of the Frame Relay circuit between R4 and R5. Configure the network so that reachability is maintained over the PPP link when R4's connection to the Frame Relay cloud is down.
- Traffic should not be routed across the PPP link when the Frame Relay circuit from R4 to R5 is up.
- Do not use the **`backup interface`** command to accomplish this.

3 Points

2.3 OSPF

- You are concerned about false routing information being injected into OSPF area 0. In order to verify the legitimacy of routing information, configure all area 0 adjacencies to be authenticated with a secure hash value of the password CISCO.

3 Points

2.4 OSPF

- Your design engineers have been performing pre-testing of new 10 Gbps Ethernet hardware for installation in your network. In order to maintain optimal bandwidth utilization throughout the OSPF domain, it is now necessary for you to manipulate how OSPF calculates its metrics.
- Configure the OSPF domain to reflect the following metric calculations:

Bandwidth (Mbps)	OSPF Cost
10,000	2
10	2000
1.544	12953
0.768	26041

2 Points

2.5 Performance Optimization

- R5's OSPF database is growing quickly, and the router spends considerable time on the database maintenance.
- Configure R5 so that LSA are grouped, checksummed, and max aged six times more often than by default.
- In order to protect R5 against flooding with the same LSA during network instability times, ensure it holds for twice the default interval before accepting the same LSA again.

2 Points

2.6 IGP Redistribution

- Redistribute where necessary to obtain full IP reachability to all advertised networks.
- R5 should route through R1 to get to the prefixes learned from BB1.
- R5 should route through R2 to get to the prefixes learned from BB3.

3 Points

2.7 BGP Filtering

- Administrators of AS 100 have been receiving complaints from users accessing resources from AS 54. After further investigation, you have determined that the majority of traffic going out towards AS 54 is transit traffic coming from AS 200 and AS 300. In order to deal with this congestion, a new corporate policy has been put into place which dictates that AS 100 cannot be used as transit to reach AS 54.
- Configure AS 100 to reflect this policy.
- This configuration should be done only on R6.

3 Points

2.8 BGP Bestpath Selection

- Advertise VLAN 3 into BGP on R3.
- AS 200 should route through AS 300 to get to these prefixes.
- This configuration should be done in AS 100.

2 Points

2.9 BGP Attribute Manipulation

- Advertise VLAN 29 into BGP on R2.
- R5 should see this prefix as follows:

```
   Network          Next Hop         Metric LocPrf Weight Path
*> 136.X.29.0/24     136.X.245.2           0             100 300 i
```

- This configuration should not affect any other prefixes on R5.

2 Points

2.10 BGP Bestpath Selection

- Administrators of AS 300 want traffic destined for VLAN 29 to come in via the PPP link between R2 and R3. Unfortunately, administrators of AS 200 have not been cooperating and have been sending all traffic for this prefix directly to AS 300 over the Frame Relay cloud.
- Configure AS 300 in such a way that all traffic destined for VLAN 29 comes in the PPP link to R3.
- In the case that this link between is down, VLAN 29 should still be accessible via the Frame Relay link.
- This configuration should be done only on R2.

3 Points

2.11 BGP AS Path

- Configure SW3 to advertise the EtherChannel link into BGP.
- Ensure R3 and SW3 will accept BGP updates with AS 100 in the AS path.
- Do not alter R2's configuration for this task.

3 Points

3. IPv6

3.1 IPv6 Addressing

- The network administrator has requested that VLAN 29 and VLAN 4 be configured to support IPv6.
- Address R2's interface Fa0/0 with the network 2001:CC1E:X:202::/64
- Address R4's interface Fa0/0 with the network 2001:CC1E:X:404::/64.
- The host portion of the IPv6 addresses should be based partly off of their interfaces' respective MAC addresses.

2 Points

3.2 IPv6 Tunneling

- Enable communication between VLAN 29 and VLAN 4 using an IPv4 based GRE tunnel.
- Use any site-local network for the IPv6 addressing within the GRE tunnel.
- Configure static routing on R2 and R4 to obtain reachability between VLAN 29 and VLAN 4.

2 Points

4. MPLS VPN

4.1 Label Exchange

- Configure label exchange between R4 and R5 using Cisco's legacy protocol.
- Make sure the TCP session does not use the Loopback0 interfaces as sources.
- Ensure reliability in case of the primary Frame-Relay interface failure.

3 Points

4.2 MPLS VPN

- Configure VLAN 77 and VLAN 44 interfaces on R5 and R4 in the VRF VPN_AB
- Use the RD value of 100:47 and two different route-target values for every VRF.
- Make sure you can ping the directly connected interfaces across the VPN cloud.

3 Points

4.3 PE-CE Routing

- Using OSPF process numbers 44 and 77 on R4 and R5 respectively, configure PE/CE routing with the respective CE devices.
- Make sure every site sees the other site's routes as inter-are summary prefixes, not an external routes.

3 Points

5. Multicast

5.1 Multicast Forwarding

- Discover the active multicast topology using the respective show commands.
- A client located on VLAN 29 has been configured to listen for the multicast group 228.22.22.22 for testing purposes. The application used to receive the multicast feed does not support IGMP.
- Configure the network so that this host can receive traffic sent to this group.
- Ensure R2 can fast switch traffic for this group out to VLAN 29.

2 Points

5.2 Multicast Filtering

- It has come to your attention that users in VLAN 4 have been abusing your Internet connection by streaming video and audio feeds during work hours. In order to prevent this unnecessary drain on your network resources, your manager has requested for you to only allow users in VLAN 4 to receive feeds for groups that are used for business related activities.
- These groups are 225.25.25.25 and 226.26.26.26.
- Configure your network to reflect this policy.

2 Points

5.3 Multicast Filtering

- Recently, you have noticed suboptimal forwarding of multicast feeds throughout your network due to problems in your unicast routing. In order to prevent multicast feeds from looping around the network, configure R1 so that it does not send any multicast traffic out its FastEthernet interface that has a TTL of less than 13.

2 Points

6. Security

6.1 Traffic Filtering

- The network administrator has requested that R6's connection to BB1 be secured to prevent unauthorized access into your network.
- Configure R6 so that it only allows TCP, UDP and ICMP traffic in from BB1 if it was originated from behind R6, or is required for another section of the lab to work.
- Ensure that users behind R6 can still traceroute to hosts beyond the Frame Relay cloud.

3 Points

6.2 DoS Prevention

- Users are complaining about slow response time to a web server at IP address 136.X.4.100. After further investigation, it appears that the web server is undergoing a HTTP SYN flood.
- In order to help deal with these attacks, configure R4 to send a TCP reset to the web server for any TCP sessions that fail to reach the established state after 15 seconds.

3 Points

7. Network Services

7.1 IOS Management

- Since some of your network administrators do not understand how to use the IOS CLI, they have requested that R4 be setup to be managed via HTTP. In order to minimize the risk of managing R4 though HTTP, use the following parameters:
    - Use TCP port 8080
    - Only permit access from the 136.X.2.0/24 subnet
    - Authenticate users using local username WEB and the password CISCO
    - This password should be stored in the router's configuration as an MD5 hash.

2 Points

7.2 File Management

- The NOC has reported that R1 has been having problems with its flash memory, and has been trying to load the default IOS image named "cisco2-C2600" via TFTP. In response to this the NOC has loaded the image"c2600-iuo-mz.122-13.bin" into R3's flash in case of a failure of R1.
- Configure the network so that R1 can boot this image from R3 if its flash fails again.

2 Points

7.3 Auto-Install

- A new router will be installed on the Frame Relay cloud connecting to R5 shortly using DLCI 555. This new router will need to get its configuration from a TFTP server located in VLAN 29.
- Configure R5 to use the 136.X.5.0/30 subnet for communication with the new router and provide it with IP address 136.X.5.2 via BOOTP.

2 Points

7.4 Local Authorization

- Following a recommendation by an outside consultant, management has requested that R2's default privilege level for telnet access be set to 0.
- The only commands (other than privilege 0 commands) that these users should be allowed to issue are ping and traceroute.
- If the users need privilege level 1 commands, they should be required to authenticate with the password CISCO prior to being given access.

2 Points

7.5 Local Authorization

- The first level support engineers from the company's NOC have complained to management that they are unable to troubleshoot RIP issues because they do not have enable access to R5. In response to this, management has decided that the NOC users should be able to turn on and disable RIP debugging, but not be allowed any other access.
- The NOC users will be entering R5 in user mode (privilege level 1).

2 Points

8. QoS

8.1 Frame Relay Traffic Shaping

- The network administrator has request that Frame Relay Traffic Shaping be configured on R1, R2, R4, and R5, according to the following requirements:
    - Data should be sent at a sustained rate of 256Kbps per DLCI.
    - In the event of congestion notification, fallback to no lower than 192Kbps.
    - Any FECNs received should be reflected as a BECN

2 Points

8.2 Rate Limiting

- In order to ensure that users on VLAN 44 are being productive during work hours, your management has requested that all HTTP responses sent out R4's interface Fa0/1 be limited to 256Kbps between the hours of 8am to 5pm Monday through Friday.
- Configure R4 to reflect this policy.

2 Points

8.3 Signaling

- Recently, you have been receiving complaints from users on VLANs 4 and 52 about low VoIP quality across the data network. After further investigation, you have determined that too much of the Frame Relay circuit between R4 and R5 is being consumed by data traffic.
- In attempt to improve VoIP performance, your network administrators have configured the client applications on these VLANs to request bandwidth reservations of the network in the transit path.
- Configure R4 and R5 to support this new setup.
- Assume that each call can reserve up to 64Kbps, and that no more than 128Kbps can be reserved at any given time.

3 Points

