IEWB-RS Volume 2 Lab 4

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	18
IPv4	6
IPv6	6
MPLS VPN	10
Multicast	8
Security	10
Network Services	14
QoS	9

GOOD LUCK!

1. Layer 2 Technologies

1.1 Traffic Control

- Enable pruning within the VTP domain.
- Although SW1 and SW3 do not have VLAN 8 locally assigned, ensure that they receive unknown unicast, broadcast, or multicast traffic for VLAN 8 over their lowest numbered trunk link to SW2.
- Traffic for VLAN 8 should not be received over any of the other trunk links.

2 Points

1.2 Spanning-Tree Protocol

- Configure SW1 as the spanning-tree root for VLAN 258.
- Configure SW3 to become the spanning-tree root for VLAN 258 in the event SW1 is not longer available.
- All VLAN 258 traffic from SW2 to SW1 should transit SW4.
- In the event that SW2's path to SW1 through SW3 is down, SW2 should use the directly connected trunk links to reach SW1 directly.
- Use the fewest number of commands to accomplish this task and do not alter SW1's port-priorities.

3 Points

1.3 Link Failure Detection

- Administrators of your network are concerned about SW1 and SW2 not being able to detect a link failure on port Fa0/15.
- Configure SW1 and SW2 so that port Fa0/15 is brought down in the case that either switch can send traffic, but not receive, or vice versa.
- As an additional precaution, configure SW1 so that interface Fa0/15 is not mistakenly elected as a designated port in the above case.

3 Points

1.4 Spanning-Tree Protocol

- Ensure the output of the **`show spanning-tree vlan 258`** command on SW3 matches the highlighted output below:

```
Rack1SW3#show spanning-tree vlan 258

VLAN0258
Spanning tree enabled protocol ieee

<output omitted>

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- ----------------
Fa0/5            Desg FWD 100       128.5    Shr
Fa0/16           Desg FWD 19        128.16   P2p
Fa0/17           Desg FWD 19        128.17   P2p
Fa0/18           Desg FWD 19        128.18   P2p
Fa0/19           Altn BLK 19        128.19   P2p
Fa0/20           Altn BLK 19        128.20   P2p
Fa0/21           Root FWD 19        128.21   P2p
```

- Do not make any changes to SW3 to accomplish this task.

2 Points

1.5 Rate-Limiting

- After monitoring the inbound utilization on SW1 Fa0/1, the network administrator has requested that SW1 be configured to limit the amount of unicast traffic received from R1.
- It has been determined that the average packet size is 954 bytes and the average number of packets is 250 per second.
- Although the average packet size is 954 bytes, the network administrator has voiced concerns that the solution should accommodate various packet sizes.
- Configure SW1 to meet these requirements using the minimal amount of commands.

3 Points

1.6 QoS

- The network administrator will be implementing QoS on SW2 in the near future and would like SW2 to be configured as follows:
    - IP Precedence 0 mapped to DSCP 0
    - IP Precedence 1 mapped to DSCP 0
    - IP Precedence 2 mapped to DSCP 0
    - IP Precedence 3 mapped to DSCP 0
    - IP Precedence 4 mapped to DSCP 32
    - IP Precedence 5 mapped to DSCP 40
    - IP Precedence 6 mapped to DSCP 0
    - IP Precedence 7 mapped to DSCP 0

3 Points

1.7 QoS

- Configure SW2 to match the output of the following command:

```
Rack1SW2#show mls qos interface fa0/2
FastEthernet0/2
trust state: trust ip-precedence
trust mode: trust ip-precedence
trust enabled flag: ena
COS override: dis
default COS: 0
DSCP Mutation Map: Default DSCP Mutation Map
Trust device: none
qos mode: port-based
```

2 Points

2. IPv4

2.1 OSPF

- Administrators of your network have reported that SW2 is low on memory. After further investigation, you have determined that a large routing table is consuming the majority of SW2's memory.
- In order to cut down on the size of SW2's routing table, configure the network so that routers in OSPF area 1 do not see any inter-area or external OSPF routes.
- Ensure that devices in area 1 maintain full reachability to the rest of your network.

3 Points

2.2 OSPF

- Configure OSPF area 2 on the Ethernet, Frame Relay, and Serial segments between R4 and R5.
- Advertise the Loopback 0 interfaces of R4 and R5 into OSPF area 2.
- You are allowed to add one additional IP subnet to accomplish this.
- Mutually redistribute between RIP and OSPF on R3 and R4.

3 Points

3. IPv6

3.1 OSPFv3

- Configure IPv6 on the Frame Relay link between R1 and R2 using the addresses 2001:141:X:12::Y/64.
- Advertise this link into OSPFv3 area 0.
- Configure IPv6 on the Ethernet link between R2, R5, and SW2 using the addresses 2001:141:X:25::Y/64.
- Advertise this segment into OSPFv3 area 1.

3 Points

3.2 OSPFv3 Summarization

- Create two Loopback interfaces in R5 and SW2 with IPv6 addresses 2001:150:X:Y::Y/64 and advertise them into OSPFv3 area 1
- Configure R1 to advertise a summary encompassing both Loopback subnets to R2.
- This summary should overlap the minimum amount of address space necessary.

3 Points

4. MPLS VPN

4.1 PE-CE Routing

- Customer routers BB1 and BB3 are in VPN named VPN_A and connect to R6 and R4..
- Enable RIP on the links connecting BB1 and BB3 to R6 and R4 respectively.
- Routes learned from BB1 that have an even third octet should be seen with a metric of 10 on R6.
- The access-list used to accomplish this should only have one line and should be effective for any additional networks learned from BB1 in the future.

3 Points

4.2 VPN Tunneling

- Allow for MPLS label exchange between R4 and R6.
- Use LDP and do not rely on dynamic routing for endpoint IP addresses discovery.
- Ensure RIPv2 prefixes are exchanged across the VPN with metric preserved.
- You are allowed to create additional interfaces to accomplish this task.

4 Points

4.3 BGP

- Enable BGP as the PE/CE protocol between R6, R4 and BB1, BB3.

- BB3 expects R4 to be in AS 100.
- Make sure you configured PE routers so that the respective CE's will not reject the advertised BGP prefixes due to AS_PATH loops.

.

3 Points

5. IP Multicast

*You have to discover the active multicast configuration using the relevant **show** commands. Do not change any of the existing configurations.*

.

5.1 Auto-RP

- Configure R2 to announce its Loopback 0 interface as a candidate rendezvous-point (RP) via Auto-RP for the multicast groups 225.0.0.0 through 225.255.255.255.
- Configure R5 to announce its Loopback 0 interface as a candidate rendezvous-point (RP) via Auto-RP for the multicast groups 239.0.0.0 through 239.255.255.255.
- SW2 should be responsible for the group to RP mappings.

3 Points

5.2 Multicast Testing

- Configure R3's interface Fa0/0 as a member of the multicast group 225.25.25.25 and interface Fa0/1 as a member of 239.39.39.39.
- Ensure that R3 responds to pings sent to these multicast groups from VLANs 12 and 43.

2 Point

5.3 Multicast Rate Limiting

- Engineers in your NOC have reported that an excessive amount of multicast traffic is being received on R3's interface Ethernet0/0. After further investigation, you have discovered that your development engineers have been testing a new multicast application on VLAN 7. Apparently these engineers have not perfected the application and it is erroneously generating multicast streams at line rate.

- In order to reduce the impact of this application on your network, configure SW1 so that no more than 1Mbps of multicast traffic is sent out towards R3.

3 Points

6. Security

6.1 Traffic Filtering

- Configure a filtering policy inbound on R6's connection to VLAN6 to conform to the following requirements:
  - Permit ICMP packets across the firewall (either direction)
  - Permit telnet to servers in VLAN 7 and VLAN 77
  - Permit HTTP and SSL access to a Web server at 141.X.88.100
  - Permit any TCP and UDP sessions initiated from behind R6 to return
  - Limit the aggregate rate of DNS and ICMP packets inbound to 128Kbps
- Use the Zone Based Firewall syntax to accomplish this task and apply the most secure inspection rules where possible.

4 Points

6.2 Spoof Protection

- Configure R4's VLAN43 interface to protect against spoofed IP packets..
- Do not use any access-lists to accomplish this task.

3 Points

6.3 Infrastructure Security

- Recent network attacks from behind BB1 prompted you to think of improving the border router security.
- In order to protect against flooding attacks targeted at R6, configure the router to limit the average rate of packets going to router's CPU to 1000 per second.
- Ensure that your configuration does not affect BGP peering sessions and router management via SSH and Telnet.

3 Points

7. Network Services

7.1 SNMP

- Configure R3 and R6 to be managed via SNMP.
- The first network management server's IP address is 141.X.7.100 and second network management server's IP address is 141.X.77.100.
- Both network management servers will be expecting the RO community string to be CISCORO, the RW community string to be CISCORW, and the community string CISCO to be used for traps.
- The first network management server will be using SNMPv1 and the second SNMP server will be using SNMPv2c.
- R3 and R6 should generate SNMP traps for changes relating to HSRP status, but these traps should only be sent to the second network management server.

3 Points

7.2 IOS Menu

- The first level support engineers from the company's NOC need to have access to R2 to ping and traceroute to R5 and R6's Loopback 0 interfaces. Since these users do not have any knowledge of Cisco IOS, the network administrator has requested that a menu be configured on R2.
- This menu should enable the NOC users to ping and traceroute to R5 and R6's Loopback 0 interfaces.
- The menu should be activated whenever the user NOC logs in using the password CISCO.
- Ensure that the NOC users can exit the menu, but do not allow them to have access to the CLI when they do so.

3 Points

7.3 DNS

- One of your network administrators has added a DNS entry that allows the NOC users to telnet to R2 by name. However, this administrator has entered the entry in DNS incorrectly to resolve to 141.X.0.22.

- Without applying this IP address to any interface permit users to telnet to R2 using this DNS entry or IP address.
- Do not use NAT to accomplish this task.

2 Points

7.4 Gateway Redundancy

- Recently, hosts on VLAN 36 suffered hours of downtime due to a hardware failure on R6. This problem was not resolved until the DHCP servers were updated to assign R3 as the default gateway for this segment. In order to prevent this problem in the future, your network team has configured half of the hosts on VLAN 36 to default to R3 and the other half to default to R6.
- Configure the network so that in the event that either R3 or R6 become unavailable, hosts on this segment should still have access to the rest of the routing domain.

3 Points

7.5 Failure Message

- Configure R3 to display a "Host Failed" message of "Connection Unsuccessful" when a telnet session to R4's Loopback 0 interface fails.

3 Points

8. QoS

8.1 Congestion Avoidance

- Utilization monitoring on R1's Ethernet segment has been indicating periods of high congestion followed by periods of low utilization. After further investigation you have determined that various TCP applications throughout the network are bursting and then backing off at the same time.
- In order to prevent this behavior, configure R1 to start dropping packets with an IP precedence of routine on this link when there are at least 15 packets in the output queue.

2 Points

8.2 Congestion Management

- Users in VLAN 45 have been complaining that it is taking a very long time to send e-mail messages through their SMTP server located in VLAN 258. After further investigation, you have determined that large file downloads from an FTP server are to blame.
- In order to reduce this slow response time configure R5 so that all SMTP packets are guaranteed at least 1.5Mbps of the output queue on VLAN258 interface.
- Do not use an access-list to accomplish this task.

3 Points

8.3 Rate Limiting

- As an additional measure to decrease response time throughout your network, configure R5 so that packets over 1250 bytes are limited to 2.5Mbps outbound on its connection to BB3.

2 Points

8.4 Link Efficiency

- Configure R4 and R5 to maximize efficiency on the PPP link by guessing character streams in frames sent over the link.

2 Points