IEWB-RS-VOL2 Lab 2

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	8
IPv4	21
IPv6	3
MPLS VPN	4
Multicast	4
Security	12
Network Services	15
QoS	12

GOOD LUCK!

1. Bridging & Switching

1.1. Link Aggregation

- Configure interfaces Fa0/19 – Fa0/20 on SW1 and interfaces Fa0/13 – Fa0/14 on SW4 to be bound as one logical layer 2 link.
- SW1 should actively negotiate this link through LACP.
- SW4 should respond to SW1's LACP requests, but should not initiate negotiation.
- Ensure that SW1 is the 801.2ad *decision maker* for this logical link.
- Both switches should actively negotiate ISL as the trunking encapsulation for this logical link.

3 Points

1.2. 802.1x Authentication

- Your network administrator has voiced some concerns relating to the security of SW1's ports Fa0/9 and Fa0/10 which are being used in the company's public meeting room. In order to provide added security for these connections, a new corporate policy mandates that clients accessing the network through these ports must authenticate prior to being granted access to the network.
- A RADIUS server with the IP address 204.12.X.100 and the key CISCO has been configured for SW1 to authenticate these clients.
- The radius server is expecting the source of these packets to come from 150.X.7.7.
- Do not use authentication on any other lines of SW1.

3 Points

1.3. Performance Optimization

- Your company has recently acquired another company and will be merging your network with the other company's network in the near future.
- The other company's network currently contains 4,000 OSPF routes.
- To ensure problems do not arise when the two networks are merged, configure SW1 and SW2 so that their routing tables can support the 4,000

routes from the new company's network plus the existing unicast routes that will exist in your network.

2 Points

2. IPv4

2.1. OSPF

- Ensure there is no DR or BDR election on the segment between R1, R2, R3 and R4.
- Do not use the **`neighbor`** statement to accomplish this.
- Authenticate the adjacency between R1 and SW1 using the clear-text password CISCO.

3 Points

2.2. EIGRP

- Enable EIGRP on VLAN 26 between R2 and R6.
- As an added security measure, configure the network so that hosts running EIGRP on VLAN 26 cannot intercept the EIGRP communication between R2 and R6.
- Configure R6 so that the rest of the routers in the EIGRP domain have only one route to the prefixes learned from BB1 with a first octet of 200.
- This route should not unnecessarily overlap IP address space.

3 Points

2.3. RIP Filtering

- Configure SW1 so that it does not accept routes with an even second octet from BB3.
- The access-list used to accomplish this should not contain more than one line.
- Do not use the **`distribute-list`** or **`distance`** keywords to accomplish this.

2 Points

2.4. IGP Redistribution

- Redistribute between RIP and OSPF on SW1.
- Redistribute between OSPF and EIGRP on R2, R3, and R4.
- Ensure that full reachability is maintained throughout the IGP domain when the Frame Relay circuit between R3 and R5 is down.

3 Points

2.5. BGP Peering

- Configure a BGP peering between R5 and BB2.
- As you are concerned about false routing information being injected from unauthorized sources configure R5 to authenticate its BGP peering session with BB2 using the password CISCO.
- Configure a BGP peering between SW1 and BB3.
- AS 400 has recently acquired from AS 100, but due to AS 54's change control policy BB3's configuration will not be modified until AS 54's normally scheduled maintenance window. As an interim solution, configure AS 400 so that BB3 still thinks SW1 is still in AS 100.

3 Points

2.6. BGP Filtering

- Part of the acquisition agreement between AS 400 and AS 100 stipulates that AS 400 will not provide transit for traffic coming from AS 54 and its customers that is destined for AS 254. Configure AS 400 to reflect this policy.

2 Points

2.7. BGP Summarization

- In order to facilitate in keeping the global BGP table as small as possible your BGP routing policy dictates that R5 should advertise one route representing your entire major network 132.X.0.0 to BB2.
- Furthermore, since the Ethernet segment between R5 and BB2 is AS 254's only connection to your network, AS 254 does not need to have any longer matches than the summary route.
- Configure your network to reflect this policy, but do not allow any other routers within your network to see the summary route.
- Do not apply any access-list or prefix-list filtering towards AS 254 to accomplish this.

3 Points

2.8. BGP Tuning

- In order to increase BGP routing stability, configure R5 to respond to IGP events with 15 seconds delay.
- At the same time, to improve BGP convergence, configure R5 to batch and send routing updates to BB2 every 3 seconds.

2 Points

3. IPv6

3.1 IPv6 Deployment

- A new corporate directive has mandated that IPv6 be deployed company-wide within the next six months. However your network administrator is concerned that when IPv6 is deployed over the WAN there may be problems with running it over the Frame Relay full-mesh network between R1, R2, R3, and R4. Therefore the administrator has requested for you to configure IPv6 on the Frame Relay connection between R2 and R3 to determine if there will be any issues with the deployment.
- Configure R2 and R3 with the addresses 2001:CC1E:X::Y/128 under their respective Loopback 0 interfaces.
- Use 2001:CC1E:X:2323::Y/64 for the Frame Relay network.

- A static route pointing to each other's /128 Loopback addresses is permitted on R2 and R3 to test reachability.

3 Points

4. MPLS VPN

4.1 L2 VPN

- You are planning to deploy a layer-2 VPN connecting circuits at R4 and R6.
- Configure R4 and R6 to provide transparent Layer-2 VPN services between VLAN6 and VLAN4 attached circuits respectively.
- Use AToM as the encapsulation method to accomplish this task.
- You are allowed to use two static routes and create additional interfaces to accomplish this.

4 Points

5. Multicast

5.1. Multicast Testing

- Discover the active multicast topology using the respective show commands.
- Configure R2's most reliable interface as the rendezvous-point (RP) for all multicast groups.
- In order to facilitate in testing reachability throughout the multicast domain your configure SW1's VLAN interface participate in the multicast group 228.28.28.28.
- Ensure that SW1 responds to ICMP echo-requests sent from VLAN 26 to this multicast group address.

2 Points

5.2. Multicast Traffic Control

- During the testing phase of your multicast deployment one of your network administrators has reported that R3 has been receiving traffic source from

VLAN 6 destined for the multicast group 228.28.28.28 even though it does not have any attached members.

- After further investigation you have found that R1 is also receiving traffic from R2 for feeds destined for R3. Configure R2 to resolve this problem.

2 Points

6. Security

6.1. Router Hardening

- After returning from a network security class one of your network administrators has convinced your manager that R5 is open to a variety of security vulnerabilities. To say the least, your manager is not happy that these vulnerabilities have been left unchecked for so long.
- In order to appease your manager configure R5 to conform to the following security recommendations:
    - Drop all source routed packets
    - Disable CDP and proxy ARP on the Ethernet segment to BB2
    - Disable BOOTP server
- A banner message should be displayed to all users that telnet into the router that states: “Access to this device or the attached networks is prohibited without express written permission. Violators will be shot on sight”.

2 Points

6.2. Zone-Based Firewall

- Your NOC engineers have noticed that R2 and R4 are being polled via SNMP from an unauthorized source.
- To avoid any problems associated with unauthorized polling configure your network so that all SNMP requests received from BB1 and BB2 are filtered out by R5 and R6 respectively.
- Use Zone-Based Firewall feature to accomplish this.

4 Points

6.3. Traffic Logging

- After blocking SNMP from outside the network it appears that the device polling R2 and R4 is internal. In order to help track down the source of the SNMP packets configure R2 and R4 to generate a log message whenever a device attempts to poll them using the Read-Only community string of ‘public’.
- These log messages should be sent to the syslog server at 132.X.33.100.

3 Points

6.4. ICMP Filtering

- Recently your network team thwarted an attempted smurf attack issued by a disgruntled ex-administrator. In order to prevent this type of attack in the future you have decided to limit the amount of ICMP traffic that R5 permits inbound on its interface attached to VLAN 52.
- Configure your network so that ICMP traffic is only allowed into your network via VLAN 52 if the traffic was initiated from behind R5.
- For diagnostic and troubleshooting purposes ensure that users throughout your network are still able to traceroute from behind R5.

3 Points

7. Network Services

7.1 RMON

- After the network was overwhelmed by the latest Microsoft® worm your network administrators have requested that R5 and R6 be configured to generate an SNMP trap whenever their average five minute CPU utilization (lsystem.58.0) reaches 75%.
- The sampling interval should be done once per minute.
- When the 75% threshold is breached, an event should be generated that reads “Five Minute CPU Average Above 75%”.
- When the utilization falls back below 40%, an event should be generated that reads “Five Minute CPU Average Below 40%”.
- The SNMP server to send these traps to is 132.X.33.100.
- The SNMP server is expecting the community string to be IETRAP.

3 Points

7.2. Remote Access

- Your manager has requested that R4 be configured to allow the users from the company's NOC to telnet in to manage the router.
- The users will expect the username to be NOC and the password to be CISCO.
- All telnet sessions should be disconnected from the router after 5 minutes of inactivity.
- The maximum amount of time a user should be allowed to telnet into R4 before being disconnected should be 15 minutes.
- Sixty seconds prior to automatically logging this user off R4 should send the user a warning message in order to give the user time to finish up and save any changes to the configuration.

3 Points

7.3. Remote Access Security

- In order to increase the security of your password database configure R4 so that the password for the NOC username is stored as an MD5 hash that represents the password CISCO.

2 Points

7.4. Syslog

- Configure R3 to log all severity 7 and below messages to a syslog server with an IP address of 132.X.33.100.
- R3 should not generate a log when its interface Serial1/0 changes status, but should generate a log when a Frame Relay DLCI changes status.

2 Points

7.5. Traffic Accounting

- Administrators of your network would like to collect usage statistics on packets that violate R5 and R6's traffic filtering policy.
- Configure R5 and R6 to collect these statistics and store them locally.
- R5 and R6 should store up to 2500 of these entries in their memory.

3 Points

7.6 System Management

- Your security manager is concerned with the possibility of unauthorized personnel gaining physical access to the networking equipment.
- In order to reduce the risk of a direct physical device exposure, ensure that no one could reload your switches and start the initial configuration using the "Mode" button on the front panel.

2 Points

8. QoS

8.1. Congestion Management

- Users behind BB2 accessing an SMTP server at 132.X.3.100 have been complaining about slow response time. After further investigation, you notice congestion on the Frame Relay link between R3 and R5.
- Configure R3 and R5 so that SMTP packets to and from the server are guaranteed at least 256Kbps during times of congestion across the Frame Relay link.
- Assume that both R3 and R5 both have a port speed of 512Kbps.

3 Points

8.2. Policy Routing

- Users in VLAN 26 have been complaining about slow response time to an FTP server located in VLAN 33 with the IP address 132.X.33.33.

- In order to alleviate congestion to and from this server a new traffic engineering policy dictates that all FTP traffic coming from VLAN 26 destined for this server and back should use the Serial link between R2 and R3 as opposed to the Frame Relay link.
- All other traffic destined to this server should follow normal forwarding.
- Assume that this FTP server does not support PASV FTP connections.
- You are allowed to use policy routing to accomplish this.

3 Points

8.3. Congestion Management

- As an additional measure to reduce FTP response time ensure that bidirectional FTP traffic between VLAN 6 and the FTP server is guaranteed a least 256Kbps on the Serial link between R2 and R3 during times of congestion.

3 Points

8.4. Frame Relay Traffic Shaping

- You have noticed a large number of frames arriving on R2's DLCI 204 and R4's DLCI 402 with the Frame Relay DE bit set. Since voice traffic will be transferred across these DLCIs in the near future your corporate policy now mandates that R2 and R4 must conform to the provider's CIR.
- Configure Frame Relay Traffic Shaping on R2 and R4 in order to resolve this issue using the following information:
    - R2's connection to the Frame Relay cloud has a port speed of 512Kbps
    - R2 has had a CIR of 128Kbps provisioned for DLCI 204 by the telco
    - R4's connection to the Frame Relay cloud has a port speed of 512Kbps
    - R4 has had a CIR of 128Kbps provisioned for DLCI 402 by the telco
- Under no circumstances should R2 or R4 burst above the provisioned CIR for these DLCIs.

- To ensure that you are not oversubscribing your Frame Relay circuits limit all other DLCIs on R2 and R4 to a rate equal to half of the remaining bandwidth (port speed minus the provisioned CIR) of the circuit between R2 and R4.
- Use the lowest interval (Tc) available for the DLCIs between R2 and R4.
- All other DLCIs should use the default interval (Tc).

3 Points