IEWB-RS-VOL2 Lab 5

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	24
IPv6	10
MPLS VPN	0
Multicast	7
Security	6
Network Services	12
QoS	11

GOOD LUCK!

1. Layer 2 Technologies

1.1 EtherChannel

- Configure an EtherChannel link between SW1's interfaces Fa0/13 and Fa0/14 and SW2's interfaces Fa0/13 and Fa0/14. Use port channel number 12.
- Configure an EtherChannel link between SW1's interfaces Fa0/16 and Fa0/17 and SW3's interfaces Fa0/13 and Fa0/14. Use port channel number 13.
- Configure an EtherChannel link between SW1's interfaces Fa0/19 and Fa0/20 and SW4's interfaces Fa0/13 and Fa0/14. Use port channel number 14.
- Do not run PAgP or LACP on these links.
- All traffic sent over these trunk links should be tagged with a VLAN header.
- Do not issue any global configuration commands to accomplish this task.

3 Points

1.2 Load Distribution

- You have noticed very high utilization on the interface Fa0/13 between SW1 and SW2 and have determined that the majority of the traffic transiting this link is coming from a single file server located behind BB2.
- Traffic is sourced from multiple clients behind R1 and R6.
- Configure the network in such a way that traffic sent over this EtherChannel link is distributed more evenly while taking into account the single server and multiple clients.

3 Points

1.3 CAM Table Maintenance

- Administrators of your network have noticed that some traffic has been leaking between VLAN 8 and VLAN 88. After further investigation, you have determined that SW2's CAM table is maxed out and has been treating some unicast frames like broadcast frames.
- In order to reduce the amount of entries in the CAM table, configure the network so that SW2 discards inactive entries from VLAN 8 and VLAN 88 after 10 seconds.

3 Points

2. IPv4

2.1 OSPF

- Configure OSPF area 27 on the Ethernet segment between R2 and SW1.
- Advertise SW1's interface Loopback 0 into OSPF area 27.
- Since SW1's only connection to the rest of the routing domain is through R2, SW1 does not need specific routing information about the rest of the network. Configure your network so that the only OSPF route that SW1 sees is a default route generated by R2.

3 Points

2.2 EIGRP

- One of the deciding factors in choosing EIGRP as an IGP for your network was the granularity of its metric calculation.
- In order to get the maximum benefit of this granularity, configure the EIGRP domain so that bandwidth, delay, and load are taken into account when computing metrics.
- Also, to ensure that bandwidth is always the major factor in metric calculation, configure the EIGRP domain so that bandwidth is three times more significant than either load or delay in the calculation.

3 Points

2.3 Routing Redundancy

- Your network administrators are concerned with a failure of R5's Frame Relay link isolating it from the rest of the network. In order to prevent this case, an additional point-to-point Serial link has been provisioned to R4.
- Configure the network in such a way that connectivity is maintained throughout the network if R5 loses its connection to the Frame Relay cloud.
- You are allowed to use static routes to accomplish this task.

3 Points

2.4 RIPv2

- Configure RIPv2 on R1, R4, and R6.
- Enable RIP on VLAN 4, VLAN 162, and the Frame Relay connection to BB1.
- Enable RIP on R6's interface Loopback 0.
- Administrators of your network are concerned about false routing information being injected into the RIP domain from VLAN 162. In order to prevent this, configure R1 and R6 to authenticate all RIP updates received on VLAN 162 with a secure hash value of the password CISCO. Use key 1 for this authentication.
- As an additional security precaution, configure R1 and R6 so that no unauthorized devices can receive RIP updates sent out on VLAN 162.

3 Points

2.5 IGP Redistribution

- Redistribute in the minimum number of places necessary to gain full reachability throughout the network.
- Routers in the OSPF domain should have the minimum amount of routes needed to reach the RIP routes learned from BB3. Do not overlap any address space to accomplish this.
- Configure summaries for the major network of your topology and for the loopback networks to advertise to BB3.
- **Note:** Networks that have not been added to an IGP (such as the networks interconnecting the switches) are not required to be reachable from the other devices in your topology.

3 Points

2.6 AS_PATH

- Create a new Loopback interface on SW1 with the IP address 162.X.7.7/24 and advertise it into BGP.
- Create a new Loopback interface on SW2 with the IP address 162.X.18.8/24 and advertise it into BGP.
- Since SW1, SW2, SW3, and SW4 have connections only to AS 300, it has been decided that they will not apply for their own block of IP addresses, nor will they apply for a public BGP AS number. Instead, AS 300 has assigned them the locally significant AS numbers of 65001, 65002, and 65034.
- Configure your network so these AS numbers do not leak out onto the rest of the network when AS 300 is advertising prefixes that have been originated in either AS 65001, AS 65002, or AS 65034.
- R1 through R6 should all be able to ping these two loopback networks.

3 Points

2.7 BGP Filtering

- Configure a new Loopback interface on R5 with the IP address 162.X.15.5/24 and advertise it into BGP.
- R4 should not pass this prefix on to any BGP speaking neighbors.
- All of this configuration should be done on R5.

3 Points

2.8 BGP Table Stability

- Recent network monitoring has shown excessive route recalculation throughout the BGP domain. After further investigation, you have found that AS 54's uplinks to the rest of the Internet have become unstable, and routes are constantly being added and withdrawn from their advertisements.
- To minimize the impact on the rest of the network, configure R4 to add a penalty of 1000 to BGP prefixes each time a withdrawn message is received for them.
- R4 should stop advertising these unstable prefixes when their penalty value exceeds 3000.
- Once a stable prefix's penalty falls below 1000, it should be reinstalled in the BGP table as an active prefix.
- Ensure that no stable prefix's advertisement is withdrawn for more than 30 minutes.

3 Points

3. IPv6

3.1 IPv6 over Frame Relay

- Configure IPv6 on the Frame Relay link between R1 and R3 using the global unicast network 2001:CC1E:X:0::Y/64.
- Configure IPv6 on the Frame Relay links between R2, R3, and R4 using the site-local network FEC0:234::Y/64.

3 Points

3.2 IPv6 BGP

- Configure IPv6 BGP peering sessions between the following devices:

Device 1	Device 2
R4	R3
R3	R2
R3	R1

2 Points

3.3 IPv6 BGP Advertisements

- Configure R1, R2, and R4 to advertise IPv6 networks of VLANs 162, 27, and 4 into BGP respectively.
- Configure R3 to advertise the IPv6 Frame Relay segments and VLAN 3 into BGP.

2 Points

3.4 IPv6 BGP Summarization

- Configure R3 so that R4 only sees one route to VLANs 3, 27, 162, and the Frame Relay link between R1 and R3.
- The advertisement should be as specific as possible while still encompassing all of these segments.
- R1 and R2 should not be affected by this configuration.

3 Points

4. MPLS VPN

No tasks in this section.

5. Multicast

5.1 RP Assignment

- Configure R3 to announce its Loopback 0 interface as a candidate rendezvous-point (RP) through Auto-RP.
- Configure R5 to announce its Loopback 0 interface as a candidate rendezvous-point (RP) through Auto-RP.
- For ease of management and future multicast configuration changes, configure R1 to map multicast groups 239.0.0.0 – 239.255.255.255 to R3 and multicast groups 226.0.0.0 – 238.255.255.255 to R5.
- Use the minimum number of access-lists and access-list entries on R1 to accomplish this.

3 Points

5.2 Multicast Features

- For security reasons, do not allow BB2 to become a PIM neighbor with R1.
- Configure your network so that SW2 will not receive traffic for any administratively scoped multicast groups regardless of any IGMP join messages it receives for these groups.
- Configure the network so that multicast groups which use R3 as their RP never change to a shortest path source tree. Instead these multicast groups should always use a shared tree.

4 Points

6. Security

6.1 Traffic Filtering

- A new corporate policy has been put in to effect that requires R4 to secure its connection to BB3. R4 should treat its interface connecting to BB3 as an 'outside' interface and all other links as 'inside' interfaces.
- Any ICMP, UDP, or TCP traffic coming in from an inside interface and exiting the outside interface should be allowed to return.
- R4 should still allow all necessary routing protocol traffic in from the outside interface but deny any other traffic to the router.
- For management purposes, R4 will need to be able to ping and telnet to BB3.
- Do not use CBAC or ZFW to accomplish this task.

3 Points

6.2 DoS Prevention

- Recently, R1 and R6 underwent a ping DoS attack that originated from behind BB2. In response to this, your network administrator has requested you to configure R1 and R6 to not receive any ICMP echo requests sourced from the 205.90.31.0/24 network inbound on their interfaces attached to VLAN 162.
- Do not apply any configuration on either R1 or R6 to accomplish this.

3 Points

7. Network Services

7.1 SNMP

- A new network management server has been installed to manage R6. Configure R6 using the following SNMP parameters:
    - Contact: CCIE Lab R6
    - Location: San Jose, CA US
    - Chassis ID: 556-123456
    - Read-Only community: CISCORO
    - Read-Write community: CISCORW
- The management station's IP address is 192.10.X.101.
- This is the only station that should be allowed to manage R6.
- Attempts by other devices to manage R6 via SNMP should be logged.
- The network management server will be expecting SNMP traps to use a community of CISCOTRAP and be sourced from R6's Loopback 0 interface.

3 Points

7.2 Syslog

- One of your network administrators has requested that R4 and R5 be configured to log all severity 5 and below messages to a syslog server with the IP address 192.10.X.101.
- This network administrator has configured the syslog server to expect these messages to use the SYS10 facility.
- R4 and R5 should include their hostname in the syslog messages.
- All syslog messages should be sourced from R4 and R5's Loopback 0 interfaces.

3 Points

7.3 DNS

- The network administrators have requested that they should be able to telnet to the routers in your network using their DNS names as opposed to their IP addresses while working on R6. The network administrator has setup a DNS server at IP address 192.10.X.100 for R6 to point to for DNS resolution.
- Ensure that if your administrators mistype a command when working on the console the router it does not try to resolve the mistyped command via DNS.
- This configuration should not affect any other lines on R6.

3 Points

7.4 Local Authorization

- The first level support engineers from the company's NOC have complained to management about not having access to view R6's running configuration.
- To appease them, configure R6 so that these users can see only the following information in the running configuration:
    - Hostname
    - Interfaces
    - Interface encapsulations
    - IP access-lists applied to interfaces
- The NOC users must enter privilege level 2 using the password CISCO prior to being able to view the configuration.

3 Points

8. QoS

8.1 Frame Relay Traffic Shaping

- Administrators in your NOC have noticed an excessive amount of packet loss across the Frame Relay cloud between R1 and R3. After further investigation, these engineers have determined that R1 has been overwhelming the Frame Relay connection to R3.
- Configure Frame Relay Traffic Shaping on R1 in order to help resolve this issue.
- R1 has a port speed of 512Kbps.
- R1's DLCI 113 has a provisioned CIR of 256Kbps.
- R1 should send data at 384Kbps and throttle down to CIR in the event of congestion notification from the Frame Relay cloud.
- In the case that R1 has accumulated credit, it should be allowed to burst up to its port speed.
- Use an interval (Tc) of 100ms and don't use any legacy commands to accomplish this task.

2 Points

8.2 RTP Header Compression

- Configure the Frame Relay connection between R3 and R4 to support RTP header compression.
- This compression should support up to 15 connections.
- R3 should only compress RTP headers if it is receiving RTP headers that are compressed.
- R3 should not perform RTP header compression with any other routers.

2 Points

8.3 Bandwidth Limiting

- Users have been complaining about slow access to servers in VLAN 27. After further investigation, one of your network administrators has reported that the congestion appears to be caused by users accessing a Microsoft SQL server in that VLAN.
- To resolve this problem, configure your network so that Microsoft SQL traffic is limited to an average rate of 256Kbps on R2's connection to the Frame Relay cloud.
- Up to 2048 SQL packets in excess of 256Kbps should be queued up by R2 before packet loss occurs.
- Do not use an access-list to accomplish this.

3 Points

8.4 Catalyst QoS

- It has been discovered that SW1 and R6 are overwhelming R1 with IP and non IP-traffic.
- In order to alleviate this problem, configure SW1 to rate-limit traffic received on VLAN162 as following:
    - TCP traffic should be limited to 1Mbps and marked with CS0 and any packets exceeding this limit should be remarked to CS1.
    - UDP traffic should be limited to 512Kbps and any packets exceeding the limit should be dropped.
    - IPX packets received by SW1 on this VLAN should be marked with CoS value of 2 if they don't exceed 128Kbps and dropped if they do.
    - Classify IPX using the EtherType value of 0x8137.
- Do not apply the policy-map to any physical interface to accomplish this.

4 Points

