IEWB-RS-VOL2 Lab 6

Difficulty Rating (10 highest): 7

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	14
IPv4	21
IPv6	8
MPLS VPN	7
Multicast	7
Security	6
Network Services	10
QoS	8

GOOD LUCK!

1. Layer 2 Technologies

1.1 Trunking

- Using 802.1Q encapsulation configure the following trunks:
  - SW1 Fa0/13 - SW2 Fa0/13
  - SW1 Fa0/16 - SW3 Fa0/13
  - SW1 Fa0/19 - SW4 Fa0/13
- SW1 should not trunk VLANs 7, 77, and 777 with SW3 and SW4.

3 Points

1.2 Spanning-Tree

- Ensure SW1 is forwarding on all trunk links for any active VLANs.
- If a new VLAN is added to the VTP domain NET12, SW1 should forward on all trunk links for the new VLAN.

3 Points

1.3 Layer 2 Tunneling

- Configure the network so that R4 and SW2 match the output below:

```
Rack1R4#show cdp neighbors fa0/1 | include SW2
Rack1SW2      Fas 0/1         121          S I       WS-C3560-2Fas 0/18

Rack1SW2#show cdp neighbors fa0/18 | include R4
Rack1R4       Fas 0/18        134          R S I     3640       Fas 0/1
```

- Use VLAN 100 if an additional VLAN is needed.
- Recabling of the network is not required for this task.

3 Points

1.4 MAC Filtering

- Port Fa0/10 of SW2 connects to an 802.11b wireless access point. Since there are only 4 hosts which should be accessing your network through this access point, the new corporate policy dictates that traffic from other hosts should not be allowed in this port. The MAC addresses of these four hosts are as follows:

Host	MAC Address
1	0050.7014.8ef0
2	00d0.586e.b710
3	00c0.144e.07bf
4	00d0.341c.7871

- Configure SW2 so that traffic is only allowed in this port if it is sourced from one of the above MAC addresses.
- In the case that other hosts try to access this port, a syslog message should be sent to the server 191.X.7.100.

3 Points

1.5 Spanning-Tree Convergence

- The Ethernet link connecting to the wireless access point has been periodically flapping and causing the wireless users to lose access to the network. After further investigation, you have determined that when the link comes up the users are subject to a delay as the spanning-tree process moves to the forwarding state.
- In order to minimize downtime for these users, configure SW2 so that this port goes immediately into the forwarding state when it is connected.
- As a precautionary measure, ensure that if a spanning-tree BPDU is received in this interface the normal forwarding delay is reinstated.

2 Points

2. IPv4

2.1 OSPF

- Configure OSPF area 0 on the Frame Relay segment between R1, R2, and R5.
- Do not use the **`ip ospf network`** statement on any of these devices.

3 Points

2.2 OSPF Filtering

- Since SW1's only connection to the rest of the routing domain is through R2, it does not need specific routing information about the rest of your network.
- Configure the network so that R2 filters all routing advertisements to SW1 with the exception of a default route.
- Do not use a distribute-list or prefix-list to accomplish this.

3 Points

2.3 Default Routing

- R3 is the only connection between the OSPF domain and the other routing domains. In order to minimize the amount of memory necessary to maintain the routing table throughout the OSPF domain, configure your network so that all routers in the OSPF network send their traffic towards R3 if they do not have a longer match in their routing table.
- In order to prevent the unnecessary forwarding of traffic that will eventually be dropped, ensure that R3 only advertises this default route if it has an active connection to either BB2 or BB3.

3 Points

2.4 IGP Redistribution

- Redistribute VLAN 32 into RIP on R3.
- Redistribute between OSPF and RIP on R3.
- All routers in the OSPF domain should have a longer match for R6's interface Loopback 0.
- No other routes should be redistributed from RIP to OSPF.

3 Points

2.5 BGP Filtering

- Memory usage on your BGP speaking devices in getting dangerously high. After investigating the problem, you have determined that the BGP table is consuming too much memory. In order to help cut down on the memory requirements throughout the BGP domain, your design team has implemented a new filtering policy. This policy states that AS 100 will not accept any prefixes from AS 54 with a mask longer than a /20.
- Configure R6 to reflect this policy.
- The prefix-list used to accomplish this should only have one line.

2 Points

2.6 BGP Summarization

- Configure R3 to advertise a summary of your major network, 191.X.0.0/16, and your Loopback 0 addresses, 150.X.0.0/20, into BGP.
- Do not use the **`aggregate-address`** command to accomplish this.
- You are allowed to use two static routes on R3 to accomplish this.

3 Points

2.7 BGP Table Stability

- High CPU utilization has been reported on R6. After further investigation, you have discovered that the prefixes 112.0.0.0/8 and 113.0.0.0/8 from AS 54's customers have been constantly flapping and causing R6 to continuously recalculate the BGP topology.
- In order to minimize the impact of this flapping on the rest of the BGP domain, configure R6 so that these prefixes are not advertised if they are consistently unstable.
- No other prefixes should be affected by this configuration.

2 Points

3. IPv6

3.1 IPv6 Addressing

- The network administrator has requested you to configure a test deployment of IPv6 between VLAN 5 and BB2.
- Configure IPv6 on the Serial connection between R2 and R3, using the network 2001:CC1E:X:23::Y/64.
- Configure IPv6 on the Frame Relay connection between R1, R2 and R5, using the network 2001:CC1E:X:125::Y/64.
- Configure IPv6 on R5's interface connecting to VLAN 5, using the network 2001:CC1E:X:5::Y/64.

2 Points

3.2 RIPng

- Enable RIPng on all interfaces running IPv6.
- Create and advertise into RIPng three additional Loopback interfaces in R3 with the following IPv6 addresses:
    - 2001:220:20:3::1/64
    - 2001:222:22:2::1/64
    - 2001:205:90:31::1/64
- Configure R3 to originate a default route to R2 via RIPng.
- R2 should not see any of the above-mentioned IPv6 subnets
- Do not use a prefix-list and do not change metrics to accomplish this.

3 Points

3.3 EIGRPv6

- Configure the IPv6 subnet 2001:CC1E:X:45::/64 on the link between R4 and R5.
- Create new Loopback interfaces in R4 with the following IPv6 addresses.
    - 2001:CC1E:X:444::4/64
    - 2001:CC1E:X:454::4/64
    - 2001:CC1E:X:484::4/64
- Configure IPv6 EIGRP AS100 between R4 and R5 and advertise a single optimal summary route for the above prefixes to R5.
- Redistribute between RIPng and EIGRPv6 to obtain full reachability.

3 Points

4. MPLS VPN

The ISP core network consisting of R4, R5 and R6 provide VPN services to the following customers:

Customer_A: BB1
Customer_B: SW3 and SW4.

VRFs, IP addressing and some basic IGP settings for these customers have been preconfigured for you. Some parts of MPLS VPN configuration might be missing and you should fill the gap yourself.

4.1 PE-CE Routing

- R6's Frame Relay link to BB1 is configured in EIGRP AS 10.
- Administrators in your NOC have reported that R6 has been generating a "neighbor not on common subnet" log message for EIGRP. After further investigation, you have determined that a provisioning error on the part of your Frame Relay service provider is to blame.
- In order to avoid security issues with this type of problem in the future, configure R6 so that it does not accept any EIGRP packets on the Frame Relay interface except those sent from BB1.

3 Points

4.2 Backup Link

- Customer_B's sites are emulated by SW3 and SW4 use OSPF as the routing protocol. .
- Configure R4 and R5 so that the preferred path between SW3 and SW4 is across the link connecting R4 and R5.
- Do not modify the OSPF process identifiers to accomplish this task.

4 Points

5. Multicast

5.1 PIM Filtering

- A media server located on VLAN 32 will be streaming a video feed to clients located on VLAN 5.
- The network administrator has requested that the Frame Relay connection between R1 and R5 be used as sparingly as possible for multicast traffic.
- To help avoid excess multicast flooding and pruning behavior over this Frame Relay connection, R1 should not allow R5 to become a PIM neighbor. However, R5 should still allow clients on VLAN 5 to receive multicast traffic for this group.
- Configure your network to support this arrangement.

3 Points

5.2 IGMP

- The network administrator has reported that clients in VLAN 363 will be using Windows® 95, which supports only IGMP version 1.
- Configure R3 to only support clients running IGMP version 1 on this interface.

2 Points

5.3 Multicast Testing

- The network administrator is trying to troubleshoot a problem relating to the multicast group 225.25.25.25 and has requested that SW1 forward traffic for this multicast group into VLAN 7. However, the testing application he is using will not be generating IGMP join messages.
- Configure SW1 to accommodate this request, but do not allow SW1 to process switch this traffic.

2 Points

6. Security

6.1 BPDU Filtering

- Recently, your managers brought in some outside consultants to perform a security audit of your network. After the audit these consultants have reported that there are unauthorized bridges in VLAN 363 sending DECnet spanning tree BPDUs. Until the source of these BPDUs can be located, the network administrator has requested that SW3 and SW4 filter off all DECnet spanning tree BPDUs in VLAN 363.
- Configure your network to accommodate this request.

3 Points

6.2 Traffic Filtering

- Recent network monitoring has shown a number of unauthorized sources attempting to telnet into SW1. In order to protect SW1 from unauthorized access, it has been decided to configure R2 to filter telnet traffic going to SW1.
- Configure the network in such a way that hosts must first authenticate to R2 before they are allowed to telnet to SW1.
- These users should authenticate with the username TELNET and the password CISCO.
- Users logging into R2 with the username CLI and password CISCO should be granted access to R2's CLI.

3 Points

7. Network Services

7.1 SNMP

- Configure R3 to be managed via SNMP. R3 will be managed by two separate network management servers.
- The first network management server's IP address is 191.X.7.100 and second network management server's IP address is 191.X.77.100.
- The network management servers will be expecting SNMP traps to use community string CISCOTRAP.
- The network management servers will be expecting the RO community string to be CISCORO and the RW community string to be CISCORW.
- Only allow the first network management server to access the RW community string.
- Allow R3 to be reloaded via SNMP.

3 Points

7.2 RMON

- The network administrator is trying to do preventative maintenance by having R1 and R3 generate a log message whenever the utilization on the HDLC link between them exceeds twice the normal rate.
- The network administrator has determined that the average change of the input octet (ifEntry.10) value for R1 and R3's HDLC link is 40000 per minute.
- Configure R1 and R3 to generate a log message whenever this value reaches twice the average rate and again when it falls back below the average rate.
- R1 should monitor 'ifEntry.10.3' and R3 should monitor 'ifEntry.10.5'.
- The sampling interval should be every 60 seconds.
- The server to log these events to is 191.X.7.100.

3 Points

7.3 CDP

- One of your network administrators has written a custom network management application that relies on CDP to determine when a neighboring device is down, and would like to test this application on the Ethernet segment between R4 and SW2.
- For this application, waiting 60 seconds between sending CDP packets is too long. Configure R4 and SW2 to send CDP updates every 5 seconds.
- In addition to this, R4 and SW2 should discard a CDP entry if the neighbor has not sent a CDP update in over 15 seconds.
- The network administrator has also requested that all CDP packets sent by R4 include its Loopback 0 interface's IP address in the packet for identification.

2 Points

7.4 UDP Echo

- Configure SW2 to respond to UDP echoes from a network management station with the IP address 191.X.77.100.
- SW2 should not respond to packets sent to the UDP 'discard' and 'chargen' ports from this network management station.

2 Points

8. QoS

8.1 Real Time Protocol

- VoIP users connected to R4 have been complaining about poor voice quality. After further investigation, it has been determined that excessive HTTP traffic being sent over the Frame Relay connection between R3 and R4 is the likely cause.
- In order to resolve this problem, ensure that all RTP packets sent over the Frame Relay circuit between R3 and R4 are prioritized.
- Allocate 25% of the bandwidth for these RTP packets.
- This configuration should be done in such a way that it is easy to add additional QoS configuration at a later date.

3 Points

8.2 Congestion Avoidance

- Even after prioritizing RTP packets, your users are still having issues with low voice quality. In order to deal with this congestion, configure your network so that HTTP traffic is dropped prior to the interface becoming congested.
- This HTTP traffic should not be reserved any bandwidth.

3 Points

8.3 Link Optimization

- The link between R1 and R3 is provisioned at only 64Kbps.
- In order to decrease the interactive traffic latency, enable a technique that reduces TCP header overhead.
- Provide resources enough to support optimization for 32 bi-directional connections.

2 Points

