IEWB-RS-VOL2 Lab 14

Difficulty Rating (10 highest): 9

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
L2 Technologies	15
IPv4	29
IPv6	6
MPLS VPN	0
Multicast	8
Security	8
Network Services	9
QoS	4

GOOD LUCK!

1. L2 Technologies

1.1 Packet Sniffing

- Users in VLAN 1011 have been reporting slow network response time, however, you have not been able to track down the problem. In order to collect more information regarding the issue, a Mac OS X host running *tcpdump* utility has been connected to port Fa0/12 of SW1.
- Configure SW1 so that all traffic received in VLAN 1011 is redirected to this host.
- The same port is to be used by the host to send regular data packets. The IP address assigned to the host is in VLAN 5.

2 Points

1.2 Configuration Management

- In order to protect against network downtime, your operations team has implemented a new policy which dictates that the current running configuration of SW1 must be archived in flash before any changed are made. Therefore, this file can be used as a reference against any newer configurations that cause problems in the future.
- This backup of the configuration should be stored in the *archive* directory and use the name *backup.config*.
- In order to make this process as simple as possible, configure SW1 so that when your administrators run the command *backup,* it automatically runs this process for them.
- Additionally, in order to minimize downtime in the event of a software crash due to faulty configuration changes, configure SW1 to load the archived backup configuration created upon the next bootup.

2 Points

1.3 Traffic Filtering

- In an effort to increase security, the network administrator has requested that SW1's port Fa0/5 be configured to only accept traffic from the MAC address 0000.0c12.3456.
- Configure SW1 to reflect this request, but do not use the command **`switchport port-security mac-address 0000.0c12.3456`** to accomplish this.

2 Points

1.4 Spanning Tree

- Disable spanning tree for VLAN 1363 on SW3 and SW4.
- Ensure that this does not create a spanning tree loop with SW4, since SW3 and SW4 are connected using two interfaces in VLAN 1363.
- Use the minimal configuration needed on SW4's interface Fa0/20 to accomplish this task without using the **`shutdown`** command.

2 Points

1.5 Switch Features

- Configure so that all traffic from R1 or R3 to another device on VLAN 1363 passes through SW3.
- Traffic from R1 to R3 should also pass through SW3.
- Do not add any VLANs or configure any trunks to achieve this task.

3 Points

1.6 Hub-and-Spoke

- Without creating subinterfaces, configure a Frame Relay hub-and-spoke network between R1, R3, and R5 with R3 as the hub.
- Traffic from R1 destined for R5 should transit R3, and vice versa.
- Use only the DLCIs specified in the diagram.
- Do not use any dynamic layer 3 to layer 2 mappings over these Frame Relay connections.
- Do not send any redundant broadcast traffic from the spokes to the hub.

2 Points

1.7 Point-to-Point Addressing

- Configure R4 and R5's Serial interfaces to use the IP addresses 167.X.45.4/32 and 167.X.45.5/32 respectively.
- Ensure that R4 and R5 can ping each other's Serial interfaces.
- The creation of additional logical interfaces is permitted for this task.
- Do not use static routing to accomplish this.

2 Points

2. IPv4

2.1 RIP

- Configure RIPv2 on R4 on the Ethernet segment connecting to BB2.
- RIP updates received on this interface should be authenticated with a secure hash value of the password CISCO.
- Recently, you have been getting complaints from users about reachability problems to prefixes learned from BB2. After consulting with the administrators of BB2, it appears that your RIP updates are getting periodically lost when sent over VLAN 42. Coincidentally, you seem to remember a recent issue with the Catalyst switch not forwarding certain multicast packets as it should. In order to see if this is in fact the problem, configure R4 so that RIP packets are sent as a broadcast instead of a multicast as they go out to BB2.

2 Points

2.2 OSPF

- One of your concerns about this migration is sub-optimal routing due to link speeds higher than 100Mbps in your network. In response to this, the other business unit has agreed that the layer 3 EtherChannel between SW1 and SW2 should be seen with a cost of 10.
- Configure your network to reflect this policy.
- Ensure that all other link costs are automatically updated accordingly.

2 Points

2.3 EIGRP

- Do not allow BB3 to intercept EIGRP updates coming from any of the EIGRP speaking devices on VLAN 1363.
- Do not use the **`neighbor`** command to accomplish this.

3 Points

2.4 EIGRP

- Do not allow EIGRP to use more than 384Kbps of the 1.536Mbps T1 link between R4 and R5.
- Enable EIGRP on the Frame Relay network between R1, R3, & R5.
- Authenticate the EIGRP adjacency between R1 & R3 with the MD5 hashed password CISCO13.
- Authenticate the EIGRP adjacency between R3 & R5 with the MD5 hashed password CISCO35.

3 Points

2.5 EIGRP

- Recently, you have noticed some inconsistencies in the EIGRP topology of various devices throughout the network. After looking into this issue further, you have discovered that R1 is low on memory and has not been computing DUAL correctly. Until the exact cause of this problem can be located, configure R1 so that it can only be used to reach networks which are directly connected to it.
- Additionally, to reduce the effect of SIA queries, set the time that R1 would wait in active state to the minimum.

2 Points

2.6 IGP Redistribution

- Redistribute between EIGRP and RIP on R4. Send a summary to BB2 for the major subnet of your topology and one that covers the loopback networks.
- Since R5 is the only place where EIGRP and OSPF meet, there is no reason for these domains to have specific reachability information about each other. Configure R5 to generate a default route into both the OSPF and EIGRP domains.

2 Points

2.7 BGP Peering

Basic BGP settings have been pre-configured in your rack as outlined below:

- In order to reduce the amount of iBGP peering sessions that need to be maintained within AS 100, R3 has been chosen as a central point of distribution for all iBGP learned routes. Your design team has notified you that additional devices will be added to your BGP network within the near future. These devices will be assigned the Loopback 0 addresses of 150.X.9.9 and 150.X.10.10. In order to ease in the integration of these and future devices into your BGP domain, the design team has suggested that you configure all iBGP peers of R3 (R1, R4, R5, & R6) in the peer group *iBGP*. Members of this group should all share the following attributes:

```
remote-as 100
route-reflector-clients
send-community
update-source Loopback0
```

- In order to prepare for the upcoming additions to R3's iBGP peers, configure these new devices as part of the peer group, however, do not allow R3 to attempt to initiate the BGP session.

3 Points

2.8 BGP Peering

- R4 has recently been acquired from AS 200, however, its upstream peer in AS 254 has not yet updated its BGP configuration. In the meantime, configure R4 to peer with BB2 in such a way that BB2 thinks R4 is in AS 200.
- This adjacency should be authenticated with an MD5 hash value of the password CISCO.

1 Point

2.9 AS-Path Manipulation

- Soon after implementing this quick fix, you received a call from an engineer from AS 200 who stated that they have lost reachability to AS 254. After working together on the problem, you and this engineer have realized that routes you are learning from AS 254 are getting prepended with AS 200 in the AS-Path, and are subsequently getting dropped when entering AS 200 downstream. Configure your network to resolve this problem.

2 Points

2.10 BGP Bestpath Selection

- Advertise VLANs 4 and 5 into BGP.
- Traffic from AS 54 destined to these prefixes should come in the Frame Relay link between R6 and BB1.
- Traffic for these prefixes should only come in from BB3 if the link between R6 and BB1 is down.

2 Points

2.11 BGP Summarization

- Advertise the Loopback0 networks of SW1 and SW2 into BGP.
- Routers outside of AS 65078 should only see one route to reach these two prefixes.

2 Points

2.12 AS-Path Manipulation

- Since SW1 and SW2's only connection to the rest of the network is through AS 100, administrators of SW1 and SW2 have decided not to apply for their own BGP AS number. Instead, AS 100 has assigned them the locally significant AS number of 65078. However, since this is not a valid public AS number, it cannot be leaked out onto the Internet.
- Configure your network so that AS 65078 is stripped out of the AS path when updates are sent to AS 100's upstream peers.

2 Points

2.13 BGP Route Injection

- Due to AS 65078's aggregation policy, AS 100 cannot implement a detailed ingress traffic engineering policy. Despite requests from your network team for AS 65078 to stop this aggregation, they have continued to do so. In response to this, your network team has had no choice but to manually re-inject the prefixes which AS 65078 has aggregated.
- Configure your network so that traffic for SW1's loopback enters the Frame Relay link between R6 and BB1.
- Additionally, all traffic for SW2's loopback should enter the Ethernet segment between R3 and BB3.
- Ensure that all other routers throughout your domain only have the aggregate block for this address space that AS 65078 has originated.

3 Points

3. IPv6

3.1 IPv6 Addressing

- Create two new Loopback interfaces on R6 with the IPv6 addresses 2001:150:X:26::6/64 and 2001:150:X:2E::6/64.
- Enable IPv6 routing in R6 and ensure that EIGRP is running on both new Loopback interface.

2 Point

3.2 IPv6 Tunneling

- Configure an IPv6 over IPv4 tunnel between R4 and R6.
- This tunnel should remain up if R4 loses the connection to either R3 or R5.
- Use the network 2001:167:X:46::/64 for this segment.

2 Points

3.3 EIGRP

- Configure the network 2001:167:X:4::/64 on R4's connection to VLAN 4.
- Enable EIGRP on VLAN 4 and the tunnel between R4 and R6.
- R6 should advertise just a single summary prefix encompassing both configured IPv6 subnets to R4.
- Additionally, R6 should advertise a default IPv6 route to R4.

2 Points

5. IP Multicast

IP multicast routing has been enabled in R3, R4 and R5. PIM sparse-mode is active on VLANs 4,5 and 1363 as well as on the Frame-Relay segments between R3 & R4 and R3 & R5. Additionally, PIM Sparse-Mode runs in the Serial link between R4 and R5.

5.1 RP Assignments

- Configure R4 as the RP for all multicast groups throughout your network.
- Recently, you read of a multicast network attack in which rogue hosts were injecting false Auto-RP messages into the PIM domain. Configure your network so that R4's RP assignment cannot be preempted by any Auto-RP learned information.

2 Points

5.2 MBONE Connectivity

- Your network design team has informed you that they would like to connect to the MBONE with a DVMRP tunnel over the Internet.
- The *mrouted* host where the tunnel will terminate has an IP address of 220.20.3.192.
- This host will be expecting the tunnel to be originated from R4 with a source address of 192.10.X.4.
- Configure R4 to reflect this request.
- Ensure that R3, R4, and R5 can use DVMRP derived information for RPF checks on multicast sources.

3 Points

5.3 DVMRP Interoperability

- Multicast sources on VLANs 4 and 5 in your will be delivering multicast feeds to hosts on the MBONE.
- Configure R4 to advertise a single route for these two networks over your DVMRP tunnel to the MBONE.

3 Points

6. Security

6.1 DoS Protection

- After further investigating the slow response time to your web server, it appears that the server is undergoing a TCP SYN DoS attack. You have reported this attack to your upstream provider for them to take the appropriate action. In the meantime, configure R4 to be a proxy for all TCP sessions initiated to this server.
- R4 should send a reset for any TCP sessions that have not reach the established state after 30 seconds.
- Additionally, R4 should start closing half-open TCP sessions after they have exceeded 1000.
- Once the amount of half-open sessions has dropped below 500, R4 should stop closing them.
- Configuration for this task should not use any commands that include the word “inspect”.

2 Points

6.2 Attack Mitigation

- While researching recent security bulletins, you have discovered your ISA web server in VLAN 4 is vulnerable to an attack from packets with malformed IP options headers.
- Configure R6 to prevent this type of attack by dropping all packets it receives from BB1 containing IP options.

2 Points

6.3 Infrastructure Security

- Change the enable secret to CISCO for R4.
- Using AAA, create two new users on R4 for administration via telnet.
    - A user named OPERATOR authenticated using the password of CISCO. Allow this user to configure any HTTP server settings and view the configured HTTP server settings. These capabilities should not be applied to any other users on R4.
    - A user named ADMIN with the password CISCO who can perform any configuration.
- Console access should not be affected
- Do not use the hierarchical privilege levels model to accomplish this task.

2 Points

6.4 Traffic Monitoring

- Configure R6 for Traffic Export with the following criteria:
- Use a profile name of BB1
- For incoming traffic from BB1, sample 2% of the traffic.
- For outgoing traffic to BB1, sample 5% of the traffic.
- Traffic should be exported out the Fa0/0 interface with a destination MAC address of 60.60.60

2 Points

7. Network Services

7.1 Configuration Management

- Recently, your operations team has suggested a new policy of backing up router configurations to your internal web server. They have requested that you create a menu system on R6 as a test deployment for level 1 engineers in the NOC to backup configurations.
- These engineers will login to R6 via telnet with the username NOC and the password CISCO.
- Once they login the following menu should appear:

```
Menu for Level 1 NOC users

    1.          View Current Configuration

    2.          Backup Current Configuration

    3.          Exit

Choose your selection:
```

- The internal web server's IP address is 167.X.5.115 and will be expecting the username NOC and password CISCO to be received via SSL at port 8080.
- R6's current configuration should be saved in the directory CONFIGS and have a filename of R6_CONFIG.txt on the web server.
- Ensure that the users can view the entire running configuration when they choose the first selection.

2 Points

7.2 NAT on a Stick

- Port Fa0/14 of SW1 connects to one of your client sites. Your design team has allocated this customer the IP address 167.1.27.2/24. Users at this site are using the private IP address space 172.16.0.0/24. Since this address space is not routable throughout your network, your client's onsite administrator has requested that you configure Network Address Translation (NAT) on the border router to hide their address space. Unfortunately, the device you are using to connect to this client is a Catalyst switch, which does not support NAT. After further investigation, you have discovered that the client does have an extra router onsite. Unfortunately, this router (R2) only has one Ethernet interface. Despite this fact, your operations team has left you with the task of determining an appropriate solution.
- Configure R2 so that these hosts can access the network.

3 Points

7.3 ICMP Error Reporting

- Traffic accounting has indicated that hosts in VLAN 5 are sending traffic to destinations that R5 does not have a route to, and that R5 is constantly informing these hosts that it cannot reach the destination in question. To reduce processor load on R5, configure it so that it only generates one of hese error messages every five seconds.

2 Points

7.4 Gateway Redundancy

- SW3's default gateway is set to 204.12.X.100.
- Configure R1 to proxy for this IP address using HSRP.
- If R1's interface S0/0 is down, R6 should proxy for this IP address.
- If R6's interface S0/0 is down and R1's interface S0/0 is down, R3 should proxy for this IP address.
- Do not use the **`standby`** command with the **`track <interface>`** option on R1.

2 Points

8. QoS

8.1 Priority Queueing

- One of your company's executives has been complaining about slow network response time. After your manager promised this executive that the problem would have been fixed by the last network upgrade, concerns are growing about the future of your IT department. In order to appease this executive and save the department you have decided to prioritize all of his traffic as it exits out to BB2.
- The executive's host resides on VLAN 4, and has an IP address of 167.X.4.204.
- Configure your network so that traffic for this host has absolute priority over all other traffic as it exits out towards BB2.
- Do not use legacy priority queueing to accomplish this.

2 Points

8.2 Congestion Management

- Recently, your customers in AS 54 have mentioned that access to your public web server is very slow. After further investigation, you have discovered that there is congestion on the Frame Relay link to BB1.
- Configure the network so that traffic from the web server is guaranteed 50% of the bandwidth of the Frame Relay circuit to BB1.
- The web server's IP address is 167.X.4.119.
- Do not use a **`policy-map`** to accomplish this.

2 Points

