IEWB-RS Volume 2 Lab 100

Difficulty Rating (10 highest): 8

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	0
IPv4	28
IPv6	6
MPLS VPN	3
Multicast	9
Security	18
Network Services	9
QoS	6

GOOD LUCK!

1. Layer 2 Technologies

All Layer 2 Settings have been preconfigured for you. There are no scenarios in this section.

2. IPv4

2.1. EIGRP Load Distribution

- One of the deciding factors in choosing EIGRP as an IGP in your network was the ability to do unequal cost load balancing. Therefore your network design specification dictates that all traffic from hosts on VLAN 18 destined for hosts on VLAN 26 be load balanced across all links in the transit path.
- Assume that the Frame Relay circuit between R1 & R2 is provisioned at 256Kbps, the circuit between R2 & R3 is provisioned at 1.28Mbps, and that the PPP link between R1 & R3 is a full T1 of 1.536Mbps.
- Configure your network so that R1 distributes traffic between R3 and R2 in a ratio of 5:1 respectively.

3 Points

2.2. OSPF

- Configure OSPF area 0 on the Frame Relay links between R3, R4, and R5.
- Do not use the **`ip ospf network`** command on R3.
- Advertise VLANs 5 and 55 into OSPF on R5.
- The VLAN 5 and 55 subnets should appear as Intra-Area routes on R3.

2 Points

2.3. RIP

- Configure RIPv2 on the Frame Relay segment between R6 and BB1.
- Redistribute between RIP and EIGRP on R6.

2 Points

2.4. IGP Redistribution

- Redistribute between OSPF and EIGRP on R3 and SW2.
- Devices in the EIGRP domain should only see only one route to the Loopback 0 subnets of R4 and R5.
- This route should not overlap any additional IP address space.

3 Points

2.5. Routing Loop Prevention

- Ensure that the RIP routes redistributed on R6 are not passed from OSPF and then back into EIGRP on R3 and SW2.
- Use a method that will automatically take into account any new routes redistributed into EIGRP from RIP on R6.

3 Points

2.6. Default Routing

- Configure R3 to originate a default route into the OSPF domain.
- This route should be withdrawn if R3 loses its connections to both R1 and R2.

3 Points

2.7. BGP Summarization

- Advertise VLAN 3 into BGP on R3.
- In order to facilitate in keeping the global BGP table as small as possible, configure your network so that AS 54 and AS 254 only see one route for your entire IP Address space.
- This advertisement need not include your Loopback address space.

3 Points

2.8. BGP Default Routing

- Since VLAN 18 is SW2's only connection to the rest of the BGP domain it does not need specific forwarding information.
- Configure your network so that SW2 sends all traffic destined for the BGP domain towards R1.
- Ensure that SW2 does not learn any other unnecessary reachability information via BGP.

3 Points

2.9. BGP Filtering

- Administrators of AS 200 have reported excessively high utilization on both the Ethernet segment connecting to AS 254 and the Frame Relay segment connecting to AS 100. After further investigation you have determined that the majority of this traffic has been coming from AS 300. In response to this, a new restriction has been put into place which restricts AS 200 from being used as transit for users in AS 300.
- Configure AS 200 to reflect this policy.
- Do not use an IP access-list or a prefix-list to accomplish this.

3 Points

2.10. BGP Reachability

- Users throughout your network have been complaining about periodic reachability problems to networks throughout the BGP domain. After further investigation you have determined that these reachability problems only occur when R4 loses connectivity to the Frame Relay cloud. However, your NOC engineers have verified that the PPP link to R5 is working correctly.
- Configure your network to resolve these users' connectivity problems.

3 Points

3. IPv6

3.1. IPv6 over Frame Relay

- Configure IPv6 over the Frame Relay segments between R1 & R2 and R2 & R3 as follows:
  - Use the network 2001:164:X:12::/64 between R1 and R2
  - Use the network 2001:164:X:23::/64 between R2 and R3
  - Use the router's number for the host addresses on these segments
- Configure IPv6 on the PPP link between R1 and R3 using the network 2001:164:X:13::/64.

3 Points

3.2. OSPFv3

- Configure OSPFv3 on the Frame Relay segments between R1, R2, and R3 using area 123.
- Configure OSPFv3 on the PPP link between R1 and R3.
- Create a new Loopback100 interface on R3 with the IPv6 address 2001:150:X:3::3/64 and advertise it into RIPng.
- Ensure that R1 prefers to reach R3's Loopback100 subnet via R2.

3 Points

4 MPLS VPN

4.1 VRF

- Add an address on SW3 and SW4 on VLAN 99 with the format 99.99.99.x/24 where x is the switch number.
- Configure SW3 for a VRF named SW3TEST, and SW4 for a VRF named SW4TEST, and add the VLAN 99 interfaces to these VRFs.
- Verify that SW3 and SW4 can ping each other from within the VRFs.

3 Points

5. IP Multicast

5.1 RP Assignment

- Configure R3 as the RP for the following multicast groups:
    - 225.10.0.0 - 225.10.255.255
    - 225.26.0.0 - 255.26.255.255
    - 225.42.0.0 - 255.42.255.255
    - 225.58.0.0 - 255.58.255.255
- Use the minimum amount of access-list entries necessary to accomplish this.

3 Points

5.2 RP Assignment

- Configure R4 as the RP for the following multicast groups:
    - 226.37.0.0 - 226.37.255.255
    - 226.45.0.0 - 226.45.255.255
    - 227.37.0.0 - 227.37.255.255
    - 227.45.0.0 - 227.45.255.255
- Use the minimum amount of access-list entries necessary to accomplish this.

3 Points

5.3 IGMP

- Your company's development engineers are testing a new multicast application on VLAN 3 that utilizes IGMPv2. In order to assist in their development process they have requested that you configure R3 to poll the segment for multicast group membership every 5 seconds.
- In addition to this they have requested that R3 prune a multicast group off the interface if the application has not responded within 3 seconds of receiving a host-query message from R3.
- Lastly, to prevent the new application from interfering with the normal operation of your network, configure R3 so that traffic from the business critical multicast feed 226.37.1.1 cannot be sent to VLAN 3 or accepted from VLAN 3.

3 Points

6. Security

6.1 Traffic Filtering

- One of your network administrators would like to access a Windows 2000 server located on VLAN 7 that is running remote desktop connection. However, your security team does not want to allow this service to be open to the entire network. As an alternative solution to leaving the service open the security team has suggested that SW1 be used to authenticate users prior to allowing them to connect to the server using remote desktop.
- Configure your network so that your administrator must authenticate to SW1 using the username RDP and the password CISCO prior to using remote desktop connection.
- Once he has authenticated to SW1 he alone should be able to access the server in this manner.
- The Windows server's IP address is 164.1.7.100.
- Remote desktop connection is listening at the default TCP port of 3389.
- To avoid a hijacking of the user's active session ensure that they must re-authenticate to SW1 every 10 minutes.

3 Points

6.2 Traffic Filtering

- After implementing the above configuration you have begun to get complaints from other network administrators that they can no longer telnet into SW1 to manage it remotely.
- In order to resolve this problem configure SW1 so that the user NOC with the password CISCO can telnet to SW1 using port 3023 to get access to the command line interface.
- Telnet at port 23 should be used just for authentication of the RDP firewall exception.
- Ensure that no other ports beside 23 and 3023 are open for users to connect to SW1 for management purposes.

3 Points

6.3 Traffic Export

- You suspect that some of your internal hosts are infected by Trojan applications and are leaking sensitive information to the external networks using data channels masqueraded under legitimate DNS requests.
- In order to collect more data, you have installed an IPS in VLAN 26 with the MAC address of 1234.5678.9abc
- Configure R6 to export all DNS request packets sourced from the internal subnet 164.X.0.0/16.
- Minimize the amount of unneeded information by only exporting packets entering from the VLAN 26 connection.

3 Points

6.4 Traffic Matching

- Configure R4 for Flexible Packet Matching to block Slammer traffic coming in from BB3.

3 Points

6.5 Zone Based Firewall

- Configure R6 for Zone Based Firewall with the following criteria:
    - Configure the Serial interfaces on R2 for the "Outside" zone.
    - Configure the FastEthernet interfaces for the "Inside" zone.
    - TCP and UDP traffic passing from the Inside zone to the Outside zone should be inspected, and return traffic should be allowed.
    - ICMP traffic should be allowed to pass freely from Inside to Outside, and from Outside to Inside.

3 Points

6.6 Local AAA

- Configure R5 for local AAA for the VTY lines with the following parameters:
    - Console lines should not be affected.
    - Configure a user with the username of CISCO and password of CISCO with full access to the device.
    - Configure a user with the username of INTERN and password of INTERN with the ability to log into the device and run a minimal set of commands, including **`show clock`** and **`show interface`** as shown below.
    - This user should not be able to make any configuration changes.


```
o
Rack1R5>show ?
  clock       Display the system clock
  flash:      display information about flash: file system
  interfaces  Interface status and configuration
  parser      Show parser commands

Rack1R5>?
Exec commands:
  <1-99>      Session number to resume
  credential  load the credential info from file system
  enable      Turn on privileged commands
  exit        Exit from the EXEC
  show        Show running system information

Rack1R5>
```

3 Points

7. Network Services

7.1. NTP

- Configure R4 as an NTP master with a stratum of 2.
- SW1 should receive NTP information from R4.
- Do not use the **`ntp server`** or **`ntp peer`** commands to accomplish this task.

2 Points

7.2 DHCP

- Configure R3's interface FastEthernet0/0 to receive its IP address via DHCP.
- R3 should use the hostname ROUTER3 for DHCP, and configure for a lease time of 28 hours.

2 Points

7.3 DHCP

- Configure R3 to send a DHCP request packet to renew its FastEthernet0/0 IP address every 3 hours.
- Do not use any interface level commands for this task.

2 Points

7.4 Network Stability

- Recently is has been discovered that R5 has some hardware issues which cause the router to reload frequently.
- Reduce the impact the reloads have on your routing topology stability by configuring R5 to suppress advertisements of its connected Ethernet interfaces into IGP for 30 seconds after reload.

3 Points

8. QoS

8.1 Legacy Frame Relay Traffic Shaping

- VoIP users on VLAN 7 have been complaining about low voice quality when dialing across the data network. After further investigation, you have determined that large file transfers have been consuming a large amount of bandwidth on the Frame Relay circuit between R3 and R4.
- The Frame Relay circuits between R3 & R4 and R3 & R5 are provisioned at 256Kbps each.
- Configure your network so that none of these devices exceed the provisioned rate on the circuit.
- To decrease the serialization delay on the circuit ensure that all the shaping intervals are the smallest possible, and that a single packet cannot take more than one interval to be transmitted.

3 Points

8.2 Queueing

- Now that your WAN circuits are properly conforming to their provisioned rate VoIP traffic sent over the circuit between R3 and R4 must be given preferential treatment.
- Configure your network so that 200Kbps of VoIP traffic is always dequeued first when it is sent over the Frame Relay circuit between R3 and R4.

3 Points

