IEWB-RS-VOL2 Lab 13

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	5
IPv4	21
IPv6	5
MPLS VPN	0
Multicast	5
Security	13
Network Services	13
QoS	17

GOOD LUCK!

1. L2 Technologies

1.1 IP Telephony

- An outside consulting firm has been hired to install Cisco 7960 IP phones throughout your network. One of the consulting firm's engineers has informed you that these phones will be sending their VoIP traffic with an 802.1P priority tag. As a test install, one of these phones has been connected to SW1's interface Fa0/22.
- Use the default VLAN for all other non VoIP traffic sent out this interface.
- Configure your network to support these requirements.

2 Points

1.2 PPP

- Configure PPP encapsulation on the Serial link between R4 and R5.
- There will be a DHCP server installed within your network in the near future.
- Configure R4 to request an IP address for its Serial interface during the IPCP negotiation process.
- R5 should forward these DHCP requests on to the server which will be installed at 139.X.11.100.
- Do not use the **`ip helper-address`** command on R5 for this task.

3 Points

2. IPv4

2.1 RIP

- Configure RIPv2 on R3.
- Enable RIP on the Ethernet segment between R3 and BB2.
- In order to prevent against a denial of service attack from false routing information being injected into the RIP domain, configure R3 to authenticate all RIP updates received on VLAN 32 with a hash value of the password CISCO.

2 Points

2.2 RIP

- Configure RIPv2 on R4, R5, and SW2.
- Enable RIP between R4 & SW2 and between R5 & SW2.
- Enable RIP on the PPP link between R4 & R5.
- Advertise the Loopback 0 interfaces of these devices into RIP.
- Configure R4 to advertise the 204.12.X.0/24 subnet via RIP, but do not send or receive RIP updates on this interface.

2 Points

2.3 RIP

- Since R5 is the only connection between the OSPF and RIP domains, R4 and SW2 do not need specific reachability information about the rest of the network.
- Configure R5 to inject a default route into RIP, to provide reachability to the OSPF domain.
- You can use one static route on R5 to ensure routing stability.
- R4 should load balance traffic destined to the OSPF domain between both R5 and SW2. Do not configure any access lists to achieve this.

2 Points

2.4 RIP

- Recently, you have been getting complaints from users on VLAN 43 that certain portions of the network are periodically unreachable. Apparently, these users lose their connection to the network and then regain it about 3 to 4 minutes later. After further investigation, you have determined that this loss of reachability coincides with the failure of the Ethernet segment between R5 and SW2, and is due to the slow convergence time of RIP.
- In order to reduce the downtime of these users, configure your network so that RIP converges 10 times as fast as the default settings.
- Ensure to maintain the default timer ratio.

2 Points

2.5 OSPF

- Configure the OSPF domain in such a way that R5 uses R1 to get to VLANs 2, 6, 7, 11, and 367.
- In the case that the Frame Relay circuit between R1 and R5 is down, this traffic should be rerouted to R2.
- Do not manipulate administrative distance or OSPF costs in any way to accomplish this task.

2 Points

2.6 IGP Redistribution

- Redistribute RIP into OSPF on R5.
- Redistribute between RIP and OSPF on R3.
- BB2 should have the minimum amount of routing information necessary to reach your network.
- Do not use the **`default`** or **`ip summary-address`** commands to accomplish this.

3 Points

2.7 BGP Aggregation

BGP synchronization should remain enabled on R4 and R6. The routers are configured in the autonomous systems per the diagram. The BGP peering sessions are preconfigured as follows:

Device 1	Device 2
R4	BB3
R4	R6
R6	BB1

- Configure R4 and R6 to advertise an aggregate of your entire major network (139.X.0.0/16) to AS 54 out both the Ethernet segment to BB3 and the Frame Relay link to BB1 respectively.
- Traffic from AS 54 and its customers which is destined for VLAN 5 should come in the Ethernet link between R4 and BB3.
- All other traffic from AS 54 destined for your network should follow normal forwarding.

3 Points

2.8 BGP Traffic Engineering

- Configure the BGP network in such a way that traffic from your devices going to prefixes learned from AS 54 with an even number in the first octet exit via the Frame Relay link to BB1.
- Traffic going to prefixes learned from AS 54 with an odd number in the first octet should exit via the Ethernet link to BB3.
- Ensure that all your devices have reachability to the BGP learned prefixes in this manner.

3 Points

2.9 BGP Filtering

- Recently, engineers in your network operations center have reported a software crash of R6. After reviewing the crash dump file created by R6, it appears that the crash was due to excessive memory utilization which had something to do with the BGP process. You suspect that this crash was due to a large fluctuation in the global BGP table, and may be due to a misconfiguration of your upstream peers.
- In order to prevent against further fluctuations in the BGP table affecting your network, configure R4 and R6 so that they will not accept more that 150000 prefixes in from AS 54.
- Additionally, configure your network so that you are alerted via syslog when the amount of prefixes learned from AS 54 exceeds 135000.

2 Points

3. IPv6

IPv6 addressing has been preconfigured on R2, R3, and R6

3.1 OSPFv3

- Enable OSPFv3 on all interfaces running IPv6.
- Ensure that R6 cannot see R2's VLAN2 prefix and R2 cannot see R6's VLAN6 prefix.

3 Points

3.2 Stateless Autoconfiguration

- Configure R6 to advertise the prefix 2001:CC1E:X:6::/64 to hosts on VLAN 6 for stateless autoconfiguration.
- These announcements should be sent unsolicited every 60 seconds.
- Hosts on this segment should consider R6 unreachable if an unsolicited advertisement isn't received within three minutes.

2 Points

5. IP Multicast

Multicast routing is enabled on R2, R3 and R5. PIM dense mode is enabled on VLAN 2, 5 and 367 interfaces, as well as on the links between R2 and R3. Additionally, PIM is enabled on the Frame-Relay link between R2 and R5.

5.1 Multicast Distribution

- Your company has recently installed a new video conferencing server in VLAN 367. Clients that will need to receive the multicast feeds generated by this video server at located in VLANs 2 and 5.
- Configure the network so that when the feed is sent from VLAN 367 to VLAN 2 it uses the HDLC link between R2 and R3, but when the feed is sent from VLAN 367 to VLAN 5 it is load balanced between R1 and R2.
- Do not enable multicast on R1 to accomplish this task.

3 Points

5.2 Multicast Routing Stability

- Configure R1 and R2 to store no more than 100 multicast routing entries in their mroute tables.
- To reduce the CPU load, configure R1 and R2 to back-off RFP checks up to 1 second interval in response to routing changes.

2 Points

6. Security

6.1 Network Hardening

- Lately, you have noticed that hosts in your network are being scanned via ICMP. After tracking down the source of these scans, you have determined that they are originating from behind BB2. After many failed attempts to get the administrator of BB2 to help stop devices from scanning your network, you have decided to secure the Ethernet connection to BB2.
- Configure R3's interface Fa0/1 to reflect the following policy:

    - Deny inbound all ICMP echo (type 8) packets.
    - Deny outbound all ICMP time exceeded and port unreachable packets to stop traceroute 'replies'.
    - Silently discard packets that are denied.
    - Log all denied packets.

2 Points

6.2 DDoS Attack Defense

- Recently, you noticed that your server at the IP address 139.X.5.100 is responding extremely slow to users HTTP requests.
- Using the “netstat” command, you noticed a lot of incomplete TCP connections entries:

```
[root@server ~]# netstat –anp
…
tcp        0      0 139.X.5.100:80              118.0.0.77:62963
SYN_RECV    -
tcp        0      0 139.X.5.100:80              118.0.0.77:62962
SYN_RECV    -
…
```

- Configure R5 to watch TCP sessions in progress and reset those that do not reach established state in 10 seconds.
- Use IOS Firewall Feature Set to accomplish this task.

3 Points

6.3 CBAC Tuning

- R5 should start dropping incomplete connections when their number exceeds 100, and stop clamping when the number reaches 80.
- Start dropping incomplete connection when their rate is above 60 per minute, and until the rate is below 40 per minute
- When the number of incomplete connections exceeds 20 per host, prevent further connections to this host for 2 minutes
- Ensure that TCP connections terminate within a 2 second window

3 Points

6.4 DHCP Security

- Configure SW1 so that only R3 is allowed to supply IP addresses via DHCP to the hosts connected to SW1.
- Allow SW1 to insert information option in DHCP messages and R3 to keep this information.
- At the same time, R1 should accept BOOTREQUEST messages even with incorrect “giaddr” field.

3 Points

6.5 Management Security

- Configure R5 to only allow telnet and SSH access for management when it is sourced from the loopback0 interface addresses of R1, R2, R3, R4, SW1, or SW2.
- You may configure an ACL for this task, however, it may not contain any deny statements, and may only contain a single permit statement.
- Management traffic from other source addresses should not be allowed.

2 Points

7. Services

7.1 Management / Logging

- Recently, a network outage was traced back to problems with the BGP peering session between R6 and BB1. To minimize the impact of a similar problem in the future, a new company policy was put into place that requires R6 to notify the network management station at IP address 139.X.2.100 whenever its BGP peering session to BB1 is lost.
- The network management station will be expecting the notifications to be sent using the community of CISCOBGP.
- For R3 and R4, you have decided to deploy a syslog server in order to store the logged access-list violations. The syslog server's IP address is 139.X.5.100.
- Configure R3 and R4 to log to this server using the syslog facility local6.

2 Points

7.2 Traffic Accounting

- Your manager has expressed interest in finding out what kind of applications users in VLAN 6 are using while at the office. Configure R6 to collect information about application traffic being sent to and received from VLAN 6 and store it locally.
- This accounting should include both the total number of packets sent and received as well as a 5 minute utilization average.
- Configure R5's Fa0/1 interface to gather output traffic statistics for flows, including packet size distribution and protocol, in addition to packet / byte counts for specific source / destination address pairs.

3 Points

7.3 DHCP

- Recently, a Windows server running DHCP was installed in your network. Your server administrators have been downloading updates and service packs for the machine for the past week, but they have informed you that there are still a few terabytes worth of updates they must install. As an interim solution, these administrators have requested that you configure R1 as a DHCP server for the network.
- R1 should supply R4's Serial interface with the IP address 139.X.45.4.
- You are allowed to use a static route to accomplish this task.

3 Points

7.4 DHCP

- R1 should supply hosts in VLAN 367 with IP addresses in the range of 139.X.0.100 to 139.X.0.200.
- The default gateway for these hosts should be R6.
- If R6 is down, R3 should be the default gateway. This behavior should not rely on any client OS-specific mechanics.
- Hosts in VLAN 367 should not have to re-lease an address once they have one.
- Additionally, these hosts should use the domain name "InternetworkExpert.com".

3 Points

7.5 Logging

- Engineers in your NOC have recently received lots of complaints from various users about a general network slow down. In response to this, one of the level 1 support engineers reloaded SW1 and SW2. After the reload, the problem went away, but the syslog messages stored in the switches' buffers were lost. This resulted in making the original problem harder to track down. This engineer recommended to management that SW1 and SW2 be configured to log their syslog messages to a real syslog server. Instead, management has asked you to configure SW1 and SW2 to store all their syslog messages locally except debug messages themselves even if they reboot.

2 Points

8. QoS

8.1 Legacy QoS Support

- You have been tasked with migrating the legacy CAR configuration on R2's interface Fa0/0 to the more flexible Modular QoS CLI. R2's CAR configuration is as follows:

```
interface FastEthernet0/0
 rate-limit input access-group 100 8000 2000 2000 conform-
  action drop exceed-action drop
 !
 rate-limit input access-group 101 128000 2000 2000 conform-
  action transmit exceed-action set-prec-transmit 0
 !
 rate-limit input access-group 102 256000 4000 4000 conform-
  action transmit exceed-action set-prec-transmit 0
 !
!
access-list 100 permit icmp any any
access-list 101 permit udp any any
access-list 102 permit tcp any any
```

3 Points

8.2 Congestion Management

- Users in VLAN 11 have been complaining about slow access to certain websites on the Internet. After ignoring their complaints for as long as you could, they have gone to your manager about the problem. After being forced to investigate the issue you have discovered a high number of output drops on R5's interface S0/0/0. Configure a QoS policy on R5 so that HTTP packets returning from the Internet destined for VLAN 11 are guaranteed 80% of the CIR value (384Kbps) outbound on S0/0's DLCI 501. Configure R5 so that the subinterface receives bandwidth information from the physical interface.

3 Points

8.3 Congestion Management

- After implementing the QoS policy, some users in VLAN 11 are still complaining about slow Internet access. After reinvestigating, you have found that large file transfers between VLAN 43 and VLAN 367 are causing latency due to the high serialization delay of these larger packets. In order to reduce this problem, configure the Frame Relay connection between R1 and R5 so that the largest serialization delay of any packet is 10ms.
- R1 and R5's port speed is 512Kbps. You can oversubscribe DLCI 501 up to R5's port speed.
- This configuration should not impact R5's DLCI 502.

3 Points

8.4 Policy Routing

- In order to ensure that this latency problem is fixed once and for all, you have decided that the file transfers between VLANs 43 and 367 be rerouted across the Frame Relay network.
- Configure the appropriate routers in your network so that packets larger than 1250 bytes sourced from VLAN 43 destined for VLAN 367 and vice versa use R2 as opposed to R1 as transit.

3 Points

8.5 VoIP QoS

- After finally solving the Internet issue for users in VLAN 11, you are now receiving complaints from VoIP users on R4 making calls to users behind BB2. These users have been complaining that voice quality has suffered since you made the changes to R5. After further investigation, you have confirmed that RTP packets are experiencing higher than acceptable latency between R4 and BB2.
- To try and solve this issue, configure a QoS policy which ensures that voice traffic receives the lowest possible latency across the Frame Relay cloud.
- Voice traffic should also be fragmented when sent across the Frame Relay cloud.

3 Points

8.6 Marking

- Configure R4 so that traffic transiting R4 and leaving the Fa0/1 interface that has SW2 either as the destination or as a transit device has the precedence set to 7.
- Other traffic to hosts on VLAN 24 should not be affected. Do not configure any access lists for this step.

2 Points

