# IEWB-RS-VOL2 Lab 8

## 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.internetworkexpert.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. It is not required to have reach ability to networks that you are not asked to add to BGP or an IGP.

## Lab Do's and Don'ts:

- Do not change or add any IP addresses from the initial configuration unless otherwise specified
- 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
- Do not create any additional VLANs
- Save your configurations often

## Grading:

This practice lab consists of various sections totaling 100 points. A score of 80 points is required to achieve a passing score. 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.

Grading for this practice lab is available when configured on Internetwork Expert's racks, or the racks of Internetwork Expert's preferred vendors. See Internetwork Expert's homepage at http://www.internetworkexpert.com for more information.

## Point Values:

The point values for each section are as follows:

| Section | Point Value |
|---|---|
| Bridging & Switching | 25 |
| IP IGP Routing | 26 |
| BGP | 13 |
| IP and IOS Features | 8 |
| IP Multicast | 11 |
| QoS | 10 |
| Security | 7 |

# GOOD LUCK!

## Note:

**There are no faults in the initial configurations**

**Do not alter the commands in the initial configurations**

# 1. Bridging & Switching

## 1.1 Trunking

- Configure three 802.1q trunks between SW1's interfaces Fa0/13 through Fa0/15, and SW2's interfaces Fa0/13 through Fa0/15.
- Configure an 802.1q trunk between SW3 and R5's interface E0/1.
- Do not run Dynamic Trunking Protocol on any of these interfaces.

**2 Points**

## 1.2 802.1q

- Configure three 802.1q trunks between SW1's interfaces Fa0/16 through Fa0/18, and SW3's interfaces Fa0/13 through Fa0/15.
- These trunks should be negotiated using Dynamic Trunking Protocol.

**2 Points**

## 1.3 ISL

- Configure three ISL trunks between SW1's interfaces Fa0/19 through Fa0/21, and SW4's interfaces Fa0/13 through Fa0/15.
- Use the minimal configuration possible on SW1 to complete this task.

**2 Points**

## 1.4 Spanning-Tree Protocol

- Configure spanning-tree according to the following requirements:
    - SW1 should be the root for VLANs 3 through 7
    - SW2 should be the root for VLANs 13 through 45
    - SW3 should be the root for VLANs 52 through 67
    - SW4 should be the root for VLANs 1 and 1001
    - No switch should be the elected root based upon a lower MAC address for any of the VLANs listed above.
    - Any other VLANs should have a root elected based on the lowest MAC address.
- Use the fewest commands needed to accomplish this task.

**3 Points**

## 1.5 Layer 2 Tunneling

- Configure SW2 and SW4 to allow communication for VLAN 26 between R2 and R6.
- Do not create VLAN 26 on either SW2 or SW4 for this task.

**2 Points**

## 1.6 Spanning-Tree Protocol

- Traffic for VLANs 3 through 7 should prefer to forward over the highest numbered directly connected trunk link to SW1.
- If the highest numbered link is down, traffic for these VLANs should prefer to forward over the next highest available directly connected trunk link.
- As a last resort, traffic for these VLANs should forward over the lowest numbered directly connected trunk link.
- This configuration should be done on SW1.

**2 Points**

## 1.7 Spanning-Tree Protocol

- Configure the network so that there is only one instance of spanning-tree for each of the following sets of VLANs:

    - VLANs 3 through 7
    - VLANs 13 through 45
    - VLANs 52 through 67
    - VLANs 1 and 1001

**2 Points**

## 1.8 Etherchannel

- Using the two remaining inter-switch links between SW2 & SW4 and two inter-switch links between SW3 & SW4, create two logical layer 3 connections.
- Use the information provided in the diagram to complete this task.

**2 Points**

## 1.9 Interface Negotiation

- One of your desktop administrators has informed you that some of the Windows machines are getting the error message: *Local Area Connection, A network cable is unplugged*. You have determined that these PCs' network cards are having trouble with the auto-negotiation of speed and duplex.
- In order to resolve this problem, ensure that all ports in VLAN 3 are hard coded to 100Mbps Full-Duplex.

**3 Points**

### 1.10 Multilink PPP over Frame Relay

- Configure a Frame Relay connection between R2's interface S0/0.203 and R3's interface S1/0.
- Configure a Frame Relay connection between R2's interface S0/0.213 and R3's interface S1/1.
- In order to maximize the utilization, configure the connection between these routers so that packets are fragmented amongst both links.
- In order to ensure a secure communication over the Frame Relay cloud, configure R2 and R3 to authenticate each other using their hostnames and an MD5 hash based off the password CISCO.

**3 Points**

### 1.11 Point-to-Point

- Configure the Frame Relay connection between R6 and BB1 using Frame Relay Inverse-ARP over PVC 100 on the main Serial interface.
- R6 should not send InARP requests for IP out any other circuits assigned to its Serial interface.

**2 Points**

## 2. IP IGP Routing

### 2.1 OSPF

- Configure OSPF area 0 on the Frame Relay connections between R2 and R3, and on the Ethernet segment between R2 and R6.
- Advertise R2 and R3's interface Loopback 0 into the OSPF domain.
- Authenticate the OSPF adjacency between R2 and R6 using OSPF type 1 authentication.

**2 Points**

## 2.2 OSPF

- Configure OSPF area 38 according to the network diagram.
- Advertise SW2, SW3, and SW4’s interface Loopback 0 into the OSPF domain.
- R3 is the only connection to the rest of the OSPF network for area 38. However, an upcoming addition to your network will involve adding another connection from SW2 to the OSPF domain. In order to maintain optimal routing while minimizing the amount of forwarding information that SW2, SW3, and SW4 need to store in the routing table, configure your network so that external LSAs are not advertised into OSPF area 38.
- Ensure that devices in OSPF area 38 still have specific forwarding information about prefixes originated in other OSPF areas.

**3 Points**

## 2.3 OSPF

- Configure OSPF area 67 on VLAN 67 between R6 and SW1.
- Advertise the Loopback 0 interfaces of R6 and SW1 into area 67.
- In order to help minimize the amount of prefixes needed throughout the OSPF domain, configure your network so that routers outside of OSPF area 67 only see one route to the Loopback 0 interfaces of R6 and SW1.
- Ensure that this summary route does not overlap any other address space.

**2 Points**

## 2.4 EIGRP

- Configure EIGRP AS 1024 on R1, R3, R4, and R5.
- Configure EIGRP on the HDLC link between R1 and R3.
- Configure EIGRP on the Frame Relay network between R1, R4, & R5.
- Configure EIGRP on the Ethernet link between R4 and R5.
- Advertise the Loopback 0 interfaces of all these devices into the EIGRP domain.
- Do not send EIGRP hello packets out any other interfaces, but do not use the **`passive-interface`** command to accomplish this.

**2 Points**

## 2.5 RIP

- Configure RIPv2 between R5 and BB2.
- In order to ensure the legitimacy of all routing updates received on VLAN 52, your corporate policy dictates that any RIP packets received on this segment should be authenticated with an MD5 hash of the password CISCO.
- Configure R5 to reflect this policy using key 1 for authentication.

**2 Points**

## 2.6 IGP Redistribution

- Redistribute between RIP and EIGRP on R5.
- Redistribute between OSPF and EIGRP where needed.

**2 Points**

## 2.7 Load Distribution

- Configure the network in such a way that traffic from R4 destined to the prefixes learned from BB2 is load balanced out the Ethernet link to R5 and the Frame Relay link to R1.
- Traffic should be distributed between the Ethernet and the Frame Relay links in a ratio of 4:1.

**3 Points**

## 2.8 IPv6 Addressing

- Enable IPv6 on R1's connection to VLAN 1001 using the network FEC0:CC1E:X:X::/64.

- Enable IPv6 on R4's connection to VLAN4 using the network FEC0:CC1E:X:X::/64.
- Enable IPv6 on the Ethernet segment connecting R4 and R5 using the IPv6 prefix 2001:CC1E:X:45::/64.
- Enable IPv6 on the Frame Relay connection between R1 and R5 using the network 2001:CC1E:X:1515::/64.
- Enable RIPng on the connection between R4 and R5 and on VLAN4 interface.

**3 Points**

## 2.9 OSPFv3

- Configure OSPFv3 area 0 on the Ethernet segment of R1.
- Configure OSPFv3 area 0 on the Frame Relay segment between R1 and R5.
- Do not use the **`ipv6 ospf network`** command to accomplish this.

**2 Points**

## 2.10 IPv6 Default Routing

- Configure R5 to advertise a default route to R1 via OSPFv3.
- When R5 receives IPv6 traffic from R1 it should drop it unless it has a longer match.
- Do not use the **`default-information originate always`** command to accomplish this.
- You are allowed one static route on R5.

**3 Points**

## 2.11 IPv6 Redistribution

- Configure the minimum redistribution necessary throughout the network so that R4 has reachability to R1's VLAN1001 network.

**2 Points**

# 3. BGP

## 3.1 BGP Peering

- Configure BGP on the following devices with the following AS numbers:

| Device | BGP AS |
|---|---|
| R1 | 65145 |
| R2 | 65267 |
| R3 | 65038 |
| R4 | 65145 |
| R5 | 65145 |
| R6 | 65267 |
| SW1 | 65267 |
| SW2 | 65038 |
| BB1 | 54 |
| BB3 | 54 |

- Configure the BGP peering sessions as follows:

| Device 1 | Device 2 |
|---|---|
| R6 | BB1 |
| R6 | SW1 |
| R6 | R2 |
| R2 | R3 |
| R3 | SW2 |
| R3 | R1 |
| R1 | R4 |
| R1 | R5 |
| R5 | R4 |
| R5 | BB3 |

- In order to reduce the amount of internal BGP peering sessions, your network designers have broken down your network into three confederated ASs: 65038, 65145, and 65267. Ensure that BGP speaking devices outside of your confederation see your network as the single AS of 100.
- Ensure that R6 advertises BGP routes learned from R2 to SW1 and vice versa.

**3 Points**

## 3.2 BGP Summarization

- Configure R5 and R6 to advertise the network 174.X.0.0/16 to BB3 and BB1 respectively.
- Do not allow any other devices in your BGP network to see this prefix.
- Use one static route on R5 and R6 each to accomplish this.

**2 Points**

## 3.3 BGP Next-Hop Processing

- Configure the network in such a way that all devices throughout your network have reachability to the BGP prefixes learned from AS 54.
- Do not advertise the Frame Relay link to BB1 or the Ethernet link to BB3 into IGP or BGP to accomplish this.
- Do not use the **`next-hop-self`** command to accomplish this.

**2 Points**

## 3.4 BGP Bestpath Selection

- Advertise VLANs 3, 4, and 7 into BGP.
- Configure the network in such a way that all traffic for VLAN 4 comes in the Frame Relay link to BB1, while all traffic for VLANs 3 and 7 comes in the Ethernet link to BB3.
- Ensure that traffic can be rerouted if there is a failure of either the link to BB1 or the link to BB3.
- Other ASs beyond AS 54 should not see these specific subnets, but instead should only see the previously advertised aggregate.

**3 Points**

### 3.5 BGP Filtering

- Advertise VLAN 1001 into the BGP domain on R1.
- Devices outside of AS 65145 should not have reachability to these network.
- Do not use any access-lists or prefix-lists to accomplish this.

**3 Points**

## 4. IP Services

### 4.1 Default Gateways

- In a sloppy attempt to provide a form of redundancy a few users in VLAN 26 have their default-gateway set to point to their own IP address as opposed to R6.
- Configure R2 and R6 not support these users.

**2 Points**

### 4.2 Web Caching

- Due to the low speed of the Frame Relay circuit that R4 uses to connect to the rest of the network, a web caching engine has been installed to provide increased web browsing performance for users in VLAN 4.
- The web servers that the users are browsing are located across the Frame Relay cloud toward R1.
- Configure R4 to support this setup, but do not attempt to cache HTTP traffic between VLANs 4 and 45.

**2 Points**

### 4.3 IP SLA

- The service level agreement (SLA) between your company and AS 54 dictates that AS 54 will guarantee 99.999% uptime and a maximum latency of 20ms on the Frame Relay link between R6 and BB1.

- In order to ensure that AS 54 is fulfilling this SLA, configure R6 to poll the Loopback address 115.0.0.1 of BB1 via 1250 byte ICMP ping packets every 30 seconds.
- R6 should account for ICMP ping packets that have a delay which exceeds 25ms.

**2 Points**

### 4.4 Gateway Redundancy

- Your network administrators are concerned about a degradation of service on the Frame Relay circuit between R6 and BB1 impacting users on VLAN 26.
- In order to avoid this problem, configure the network in such a way that users in VLAN 26 use R6 as their default gateway, but only if AS 54 is honoring the service agreement on the circuit between R6 and BB1.
- If AS 54 is in violation of this agreement, users in VLAN 26 should use R2 as their default gateway.
- The network administrators have informed you that all these users have their default gateways set to 174.X.26.254.

**2 Points**

## 5. IP Multicast

### 5.1 PIM

- Configure IP Multicast routing on R1, R2, R3, R4, and R5.
- Configure PIM on the Frame Relay segments between R1, R4, & R5, and R2 & R3.
- Configure PIM on the HDLC link between R1 and R3.
- Configure PIM on VLANs 1001, 26, 3, 4, and 52 of R1 through R5 respectively.

**2 Points**

## 5.2 Auto-RP

- Configure R1 and R2 to announce their Loopback 0 interface as candidate RPs via Auto-RP.
- Configure R3 to map all multicast groups with an even numbered first octet to R1 and odd-numbered first octet to R2.

**2 Points**

## 5.3 Multicast Distribution

- Configure your network so that all multicast traffic switches over to a source based tree once the source is sending at a rate greater than or equal to 128Kbps.

**2 Points**

## 5.4 Multicast Testing

- Recently, your network administrator has reported that clients in VLAN 4 cannot receive multicast feeds from servers located in VLAN 52.
- Configure the network in such a way to resolve this problem, and so that R4 responds to ICMP echo requests sent to the multicast group 226.0.0.4 sent from VLAN 52.
- Do not use tunneling to accomplish this.

**2 Points**

## 5.5 Broadcast Distribution

- Market analysts from your finance department have had a stock ticker application installed on VLAN 26. They have requested that users on VLAN 1001 also be able to access the data generated by this application. Unfortunately, this is a proprietary application in which the server only supports sending traffic to the all subnet broadcast address (255.255.255.255) using UDP port 3434, and the client only supports receiving broadcast traffic sent to this port.

- Configure your network so that hosts in VLAN 1001 can receive this market feed.

**3 Points**

# 6. QoS

## 6.1 Frame Relay Traffic Shaping

- After reviewing the monthly utilization report from the Frame Relay service provider, you have noticed that an excessive amount of frames are being marked as Discard Eligible on the Frame Relay connections between R1, R4, and R5.
- To resolve this, configure these devices to conform to their subscribed CIRs in accordance with the Frame Relay service provider.
- Each circuit has been provisioned at 128Kbps.

**2 Points**

## 6.2 Queueing

- After configuring Frame Relay Traffic Shaping, administrators in the NOC have reported an excessive amount of output drops on R1's connection to the Frame Relay cloud.
- To resolve this, configure R1's traffic shaping queue to hold 10 times the default amount of packets.

**2 Points**

## 6.3 Congestion Management

- You have noticed that delay-sensitive audio traffic (UDP destination port of 7070) sent over the Serial link between R1 and R3 is experiencing an unacceptable amount of latency due to the high amount of data traffic that is transiting the link.

- Configure the network so that this delay sensitive audio traffic is given priority over any other traffic sent across the Serial link.
- This audio traffic should be allocated a maximum of 128000 bps of the output queue of both R1 and R3.
- Your NOC engineers have told you that this audio traffic has the tendency to be sent in short bursts. Ensure to allow for a burst value of 64000 bits.

**3 Points**

### 6.4 Congestion Avoidance

- Network monitoring has indicated an inordinate amount of output drops accumulating on the Frame Relay connection of R4. After investigation, you have discovered that this is due to traffic originating from the 100 Mbps FastEthernet segment of VLAN 4 exiting the 128 Kbps Frame Relay circuit.
- In order to prevent this type of tail drop behavior for voice traffic, configure R4 to randomly drop packets on the Frame Relay circuit before congestion occurs.
- In order to ensure that voice traffic gets better service than other traffic, configure R4 so that ‘critical’ traffic will not be dropped unless there are 60 packets in the output queue.
- If there are 90 critical packets in the output queue, R4 should randomly drop 5 out of every 25 of these packets.
- In the case that there are more than 90 critical packets in the output queue, they should all be dropped.

**3 Points**

## 7. Security

### 7.1 Router Hardening

- After returning from a network security class, one of the network administrators has convinced your manager that R5 is open to many security vulnerabilities. Your manager is not happy that these vulnerabilities have been left unchecked for so long.

- In order to appease him, configure R5 to conform to the following recommendation:

  - Drop all source routed packets
  - Disable proxy-arp and CDP support on the connections to BB2 and BB3.
  - Drop all HTTP and telnet sessions destined for the 174.X.0.0/16 and the 150.X.0.0/16 networks coming from BB2 or BB3.
  - Drop all inbound echo requests coming from BB2 or BB3.

**3 Points**

## 7.2 Traffic Filtering

- Network monitoring has indicated that BB2 and BB3 are using R5 as a transit device to get to each other. In order to avoid the liability of a network attack transiting your network between these two providers, your corporate policy dictates that all traffic coming from BB2 destined for BB3 and vice-versa should be dropped.
- Configure R5 to reflect this policy, but do not use any access-lists to accomplish this.

**2 Points**

## 7.3 Traffic Filtering

- After implementing the transit filter on R5, your manager has received a request from the administrator of BB2 to allow SMTP traffic between a server in VLAN 52 and its clients in VLAN 53.
- The SMTP server's IP address is 192.10.X.100.
- Configure R5 to reflect this policy.

**2 Points**