IEWB-RS-VOL2 Lab 7

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.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.

Lab Do's and Don'ts:

- Do not change 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
- 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
Troubleshooting	3
Bridging & Switching	26
IP IGP Routing	28
BGP	17
IP and IOS Features	8
IP Multicast	3
QoS	10
Security	5

GOOD LUCK!

Troubleshooting

- There are 3 issues in the initial configurations that need to be resolved.
- Each issue is worth 1 point.

3 Points

1. Bridging & Switching

VLANs have been created and assigned to the switchports according to the diagram supplied with the task.

1.1 Trunking

- Configure SW1 interface Fa0/16 and SW3 interface Fa0/13 as a trunk link.
- Configure SW1 interface Fa0/19 and SW4 interface Fa0/13 as a trunk link.
- These trunks should be negotiated via DTP.

2 Points

1.2 EtherChannel

- Configure SW1 and SW2 according to the highlighted output below:

```
#show etherchannel 13 port-channel
                Port-channels in the group:
                ---------------------------

Port-channel: Po13
------------

Age of the Port-channel   = 00d:00h:02m:38s
Logical slot/port   = 2/13          Number of ports = 3
GC                  = 0x00000000      HotStandBy port = null
Port state          = Port-channel Ag-Inuse
Protocol            =    -

Ports in the Port-channel:

Index   Load   Port     EC state        No of bits
------+------+------+------------------+-----------
  0     00     Fa0/13   On/FEC             0
  0     00     Fa0/14   On/FEC             0
  0     00     Fa0/15   On/FEC             0

Time since last port bundled:    00d:00h:01m:34s    Fa0/15

#show interfaces po13 trunk

Port        Mode         Encapsulation  Status        Native vlan
Po13        on           802.1q         trunking      1

Port        Vlans allowed on trunk
Po13        1-6,8-4094

Port        Vlans allowed and active in management domain
Po13        1,3-6,42,55,57,263

Port        Vlans in spanning tree forwarding state and not pruned
Po13        1,3-6,42,55,57,263
```

2 Points

1.3 802.1q Tunneling

- Configure the SW1 and SW4 in such a way that R3 Fa0/0 and SW2 Fa0/20 appear directly connected via CDP.
- R1 Fa0/0 and SW2 Fa0/21 should also appear directly connect via CDP.
- If additional VLANs are needed, use VLANs 100 and 101.

3 Points

1.4 Rate-Limiting

- Configure SW1 Fa0/5 to limit unicast traffic inbound to 25% of the interfaces bandwidth.
- Use the minimal configuration possible for this task.

2 Points

1.5 IP Telephony

- Ports Fa0/7 and Fa0/8 on SW2 connect to Cisco 7960 IP phones.
- These IP phones will need to communicate with a CallManager server that is located in VLAN 4.
- Using the *minimal* configuration possible ensure that the IP phone's VoIP traffic can communicate with the CallManager server.
- The PCs connected to the IP Phones use the same VLAN number as the voice traffic.

2 Points

1.6 IP Telephony

- VoIP originating from these IP phones is being marked with a layer 2 CoS value of 5.
- Ensure that VoIP traffic originating from the IP phones maintains its CoS value as it is processed by SW2, while traffic originating from the PCs is remarked with a CoS of 1 by the IP phone.
- Traffic coming from the PCs connected to the access ports of the IP phones should be assigned to VLAN 5.

2 Points

1.7 EtherChannel

- Using all remaining inter-switch links between SW1 & SW4 and SW3 & SW4, configure two layer three EtherChannel links.
- PAgP should negotiate these EtherChannel links unconditionally.
- Use the port-channel numbers and IP addressing specified in the diagram.

2 Points

1.8 Private VLANs

- In the near future, two new servers will be added to the network.
- The first server will be connected to SW1 port Fa0/9 and the second server will be connected to SW2 port Fa0/9.
- These servers will be using IP addresses from the 192.10.X.0/24 network.
- Configure the network to meet the following requirements:
  - Both servers should not be able to communicate with each other directly
  - Both servers should be able to communicate with R4 and BB2
- If an additional VLAN is needed, use VLAN 500

3 Points

1.9 Point-to-Point

- Configure the Frame Relay connections between R3 & R5 and R4 & R5 per the diagram.
- Do not use subinterfaces on R3 or R4.
- Do not use Frame Relay Inverse-ARP.
- Do not use the **`frame-relay map`** command on R5.

2 Points

1.10 Circuit Tracking

- The Frame Relay connection between R4 and R5 is serviced by two separate Frame Relay service providers. These providers do not inform each other about the status of their local DLCIs. This, in turn, can cause one side's DLCI to remain *active* even though the other side's interface is down.
- To detect this situation and to bring R5's subinterface to R4 down if this occurs, configure R5 to poll R4 for their Frame Relay connection status.
- R4 should be configured to respond to the polls from R5 but not initiate them.

2 Points

1.11 Point-to-Point

- Configure a Frame Relay connection between R1 & R2.
- Do not use subinterfaces on R1 or R2.
- Do not use Frame Relay Inverse-ARP.

2 Points

1.12 PPP over Frame Relay

- Due to your service provider's security policy, PPP is required over the Frame Relay circuit to BB1 for the purposes of authentication.
- Configure R6 to run PPP over PVC 201 connecting to BB1 using interface Virtual-Template1.
- BB1 will be sending an authentication challenge with the username BB1.
- R6 should reply with the username ROUTER6 and an MD5 hash value of the password CISCO.
- Do not use the global command **`username`** to accomplish this.

2 Points

2. IP IGP Routing

2.1 RIP

- Configure RIP on R1, R2, R3, R4, R5, R6, and SW2.
- Enable RIP on the following links:
  - VLAN 6 on R6.
  - The Ethernet segment between R1 and SW2.
  - The Ethernet segment between R3 and SW2.
  - VLAN 263 between R2, R6, and BB3.
  - The PPP link between R4 and R5.
  - The Frame Relay segment to BB1.
- Advertise the Loopback 0 interfaces of R6 and SW2 into RIP.
- Do not send RIP updates out any other interfaces.

2 Points

2.2 RIP Filtering

- Configure R6 so that it does not advertise its route for the Frame Relay network to either R2 or BB3.
- Configure both R2 and R6 to not accept any RIP advertisements from either BB1 or BB3.

3 Points

2.3 Default Routing

- Configure RIP between R5 and SW1.
- Advertise SW1's interface Loopback 0 into RIP.
- SW1's only connection to the rest of the routing domain is through R5, therefore it does not need specific forwarding information about any prefixes.
- Configure your network so that the only IGP route SW1 sees is a default route from R5.

2 Points

2.4 RIP

- Enable RIP on VLAN 42 on R4.
- Configure MD5 authentication on the RIP session between R4 and BB2 using key 1 and the password CISCO.
- The only route that should be advertised to BB2 through RIP is a summary of your internal address space 163.X.0.0.
- This summary should encompass your entire network and still be as specific as possible without unnecessarily overlapping address space.
- Do not accept any RIP advertisements from BB2.

2 Points

2.5 OSPF

- Configure OSPF area 0 on the Frame Relay circuit between R4 and R5.
- Configure OSPF area 1 on the Frame Relay circuit between R3 and R5.
- Use the most appropriate OSPF network type for these links.

2 Points

2.6 OSPF

- Advertise the VLANs 4, 5, and the Loopback 0 networks of R4 and R5 into OSPF area 0.
- R3 should see the route to the Loopback networks as follows:

```
R3#show ip route 150.X.5.5
Routing entry for 150.X.4.0/23
  Known via "ospf 1", distance 110, metric 782, type inter area
  Last update from 163.X.35.5 on Serial1/0, 00:00:16 ago
  Routing Descriptor Blocks:
  * 163.X.35.5, from 150.X.5.5, 00:00:16 ago, via Serial1/0
      Route metric is 782, traffic share count is 1
```

2 Points

2.7 OSPF

- Configure OSPF area 1 on the Serial link between R1 and R3.
- Configure OSPF area 2 on the Frame Relay circuit between R1 and R2.
- Do not send multicast OSPF packets over either of these links.
- R1 should be elected the Designated Router for both of these circuits.

2 Points

2.8 OSPF

- Advertise the Loopback 0 networks of R1, R2, and R3 into OSPF.
- These networks should all appear with a subnet mask of /24 throughout the OSPF domain.
- Do not use any interface level **`ip ospf`** commands to accomplish this.

2 Points

2.9 OSPF

- Configure OSPF area 0 between SW1 & SW4 and SW3 & SW4.
- Configure OSPF area 0 on SW3's interface VL3.

1 Point

2.10 OSPF

- Advertise SW3 and SW4's Loopback 0 interfaces via OSPF.
- These two networks should appear as below in SW1's routing table.

```
Rack1SW1#show ip route ospf | include _10
O IA 10.0.0.0/8 [110/2] via 163.1.0.4, 00:00:56, Port-channel14
```

- SW1 should not have any other OSPF routes for subnets within the 10.0.0.0/8 network
- Do not use redistribution to accomplish this task.

2 Points

2.11 IGP Redistribution

- Redistribute between RIP and OSPF on R1, R2, R3, R4 and SW1.
- Ensure that SW2 uses the most optimal routing path to reach all prefixes in the IGP domain. This configuration should be done on SW2.

2 Points

2.12 IPv6 Addressing

- Configure IPv6 on the Frame Relay link between R3 and R5 using the network FEC0:CC1E:X:35::/64.
- Configure IPv6 on the Frame Relay link between R4and R5 using the network FEC0:CC1E:X:54::/64.
- Configure IPv6 on the Serial link between R4 and R5 using the network FEC0:CC1E:X:45::/64.
- Enable IPv6 on VLANs 4 and the Fa0/0 interface of R3 using the networks FEC0:CC1E:X:4::/64 and FEC0:CC1E:X:38/64 respectively.

2 Points

2.13 RIPng

- Create new Loopback interfaces on R4 with the following IPv6 addresses:
  - 2001:220:20:3::1/64
  - 2001:222:22:2::1/64
  - 2001:205:90:31::1/64
- Configure RIPng on all interfaces running IPv6.
- Traffic from the subnet of R3's Fa0/0 interface destined for the prefixes learned from BB2 should use the point-to-point Serial link between R4 and R5.
- If this link is down, traffic from the subnet of R3's Fa0/0 interface to these prefixes should be rerouted over the Frame Relay circuit between R4 and R5.

2 Points

2.14 IPv6 Filtering

- Configure R4 so that hosts running IPv6 on the subnet of R3's Fa0/0 interface do not have access to IPv6 enabled hosts in VLAN 4.
- Do not use a prefix-list to accomplish this.

2 Points

3. BGP

3.1 BGP Peering

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

Device	BGP AS
R1	100
R2	100
R3	300
R6	100
SW2	300

- Configure the BGP peering sessions as follows:

Device 1	Device 2
R6	BB1
R6	BB3
R6	R2
R2	BB3
R2	R1
R1	R3
R1	SW2
R3	SW2

- Ensure that all routers inside AS 300 include the community attribute when sending updates to iBGP neighbors.

2 Points

3.2 BGP Peering

- Configure BGP on R4, R5, and SW1 using AS numbers 65004, 65005, and 65007 respectively.
- R5 should peer with R3, R4, and SW1.
- R4 should peer with BB2, and use the password CISCO for authentication.
- From the perspective of the rest of the BGP network, R4, R5, and SW1 should all appear to be members of AS 200.

3 Points

3.3 BGP Bestpath Selection

- For the purposes of load balancing and redundancy, AS 100 has multiple connections to AS 54.
- In order to more evenly distribute the traffic load, configure your network so that all traffic from AS 100 destined for prefixes originated in AS 54 transits the link to BB1.
- In addition to this, configure your network so that all traffic from AS 100 destined for prefixes that are from customers of AS 54 is sent out towards BB3.
- In the case that the link to BB1 is down, traffic for prefixes that have been originated inside AS 54 should still be able to be rerouted to BB3.
- All of this configuration should be done on R6.

3 Points

3.4 BGP Next-Hop Processing

- Since the Frame Relay link between R6 and BB1 is only used for transit, there is no reason for anyone else in the routing domain to have a route to this prefix. Therefore, in order to facilitate in keeping your network's routing table as small as possible, do not advertise this prefix into either IGP or BGP.
- Ensure that all routers throughout your network still have IP reachability to all BGP prefixes learned from AS 54.

2 Points

3.5 BGP Failure Detection

- Administrators of your network have been reporting reachability problems to prefixes originated in AS 54 and suspiciously high CPU utilization on R6. After further investigation, you have determined that R6 has been constantly recalculating the BGP topology due to the Frame Relay link flapping. In response to this problem your support team has opened a trouble ticket with the telco, but does not realistically expect a response for a few weeks. In the meantime you must minimize the amount of time R6 spends recalculating the BGP table.
- Configure your network so that if the Frame Relay circuit on R6 goes down the BGP peering session with BB1 is not declared down until a hello packet has not been heard for 30 seconds.

2 Points

3.6 BGP Aggregation

- AS 200's only path to AS 100 and its customers is through AS 300. Since this is the case, BGP speakers outside of AS 300 do not need specific forwarding information about AS 100's customers.
- In order to reduce the size of the global BGP table, configure your network so that all BGP speaking routers in AS 200 and beyond see the minimum amount of prefixes necessary to reach AS 100's customers.
- Do not use any default routing to accomplish this.
- Ensure not to overlap any address space when configuring this summarization.
- This configuration should be done on R1.

2 Points

3.7 BGP Traffic Engineering

- AS 300 also has multiple connections to AS 100. However, due to the aggregation recently configured in AS 100, AS 300 can no longer implement a detailed traffic engineering policy. In order to maximize the utilization on both links connecting AS 300 and AS 100, the administrators of both ASs have agreed on the following traffic engineering policy
- All traffic for the following destinations should transit VLAN 18:
  - 28.119.16.0/24
  - 112.0.0.0/8
  - 113.0.0.0/8
  - 114.0.0.0/8
  - 115.0.0.0/8
- All traffic for the following destinations should transit the Serial link between R1 and R3:
  - 28.119.17.0/24
  - 116.0.0.0/8
  - 117.0.0.0/8
  - 118.0.0.0/8
  - 119.0.0.0/8
- Autonomous Systems beyond AS 300 should still have the minimum amount of routes necessary to reach all prefixes learned from AS 100.
- All of this configuration should be done in AS 100.
- Configure your network to reflect this policy.

3 Points

4. IP and IOS Features

4.1 Syslog

- Management has implemented a new policy that requires all devices to log their syslog messages to 163.X.5.100 and 163.X.6.100.
- Edge routers (R2, R4 and R6) should log using facility local3.
- Internal routers (R1, R3, and R5) should log using facility local4.
- Switches (SW1, SW2, SW3, and SW4) should log using facility local5.
- These log messages should be time stamped with the current date and time, including the millisecond.

2 Points

4.2 NTP

- After implementing syslog, your NOC engineers have noticed inconsistent timestamps on the syslog messages. Therefore, they have requested for all devices to receive network time from BB3.
- BB3 has filtering in place for NTP packets and will be expecting the NTP requests to be sourced from each your devices’ Loopback 0 interfaces.

2 Points

4.3 DHCP

- The network administrator has requested that R6 respond to DHCP requests for clients in VLAN 6.
- R6 should provide clients with the following information:
  - IP addresses: 163.X.6.128 though 163.X.6.250
  - Exclude IP address: 163.X.6.130
  - Default Gateway: 163.X.6.6
  - Domain Name: InternetworkExpert.com

2 Points

4.4 Network Roaming

- Your accounting department has recently purchased a custom software package that has been specifically licensed for a PC in VLAN 5 with the IP address of 163.X.5.25. Due to new construction, your accounting department will be shortly relocated to a different portion of your building, and will therefore connect to your network through a different VLAN. However, the accounting department does not want to pay the software company a fee to have the license changed to the new IP in VLAN 6.
- Configure the network in such a way that this PC can function properly when moved to VLAN 6.
- Do not allow any other hosts to access the network in this manner.

2 Points

5. IP Multicast

5.1 PIM

- Recently, one of your users in VLAN 4 has requested access to a multicast feed from a media server located in VLAN 5.
- Configure PIM on VLANs 4, 5, and the Serial link between R4 and R5 to accommodate this user’s request.
- Do not use any rendezvous point assignments to accomplish this.

3 Points

6. QoS

Recently, users in VLANs 4 and 5 have been given access to a VoIP based application to communicate with each other over your data network. This application uses TCP port 1720 for H.323 signaling and UDP ports 16384-32767 for actual voice payload. In order to ensure that the VoIP traffic gets the expedited forwarding it requires, your administration has clearly defined a strict end-to-end QoS policy for your network. This policy will utilize DSCP values to differentiate between various data and voice traffic classes throughout your network while maintaining backwards compatibility with IP precedence values, and should be implemented as follows.

6.1 Marking

- The first step in your end-to-end QoS policy is to ensure that all traffic is properly categorized. To do so, configure all VoIP signaling and payload traffic coming from VLANs 4 and 5 to be marked with a DSCP value of CS5 for critical. All non VoIP traffic should be marked with a DSCP value of CS1 for routine.
- In order to ensure that all other data traffic does not get expedited service, configure the voice domain so that packets received on the network edge are rewritten with the appropriate DSCP value.

2 Points

6.2 Shaping

- The next portion of the QoS policy dictates that all traffic sent across the Frame Relay cloud should be shaped, in order to not to cause congestion for the VoIP traffic.
- The Frame Relay interfaces of R3, R4, and R5 are all clocked at T1 speed by the Frame Relay service provider. However, since R5 only has a single connection to the Frame Relay cloud, each VC on R5 has been equally provisioned a CIR of 768Kbps by the telco.
- Configure all endpoints of the Frame Relay network to adhere to the provisioned CIR.
- The shaping intervals between R4 and R5 should be such as to minimize delay due to the serialization of the interface.
- As an additional measure to decrease the delay of your VoIP traffic, configure R4 and R5 so that packets with a payload greater than 960 bytes are fragmented.

3 Points

6.3 Marking

- To ensure that your voice traffic is not dropped in the case that the Frame Relay cloud experiences congestion, configure your network so that all non-VoIP traffic sent across the provider cloud has the Frame Relay discard eligibility bit set.

2 Points

6.4 Prioritization

- The last portion of your QoS policy states that VoIP traffic must be given preferential treatment over other traffic classes.
- To accomplish this, configure your network so that R4 and R5 always sends VoIP traffic out the Frame Relay circuit between them and the VLAN 4 & 5 segments before any other traffic.
- In order to ensure that your other traffic classes do not get starved of bandwidth, configure your network so that if there is more than 256Kbps of VoIP traffic in the output queue and there is congestion, excess VoIP traffic is dropped.
- When there is no congestion, VoIP traffic above 256Kbps may be sent, but it should not be guaranteed low latency.

3 Points

7. Security

7.1 DoS Prevention

- The network administrator is concerned about the possibility of older Windows clients in VLAN 4 being the victim of a DoS attack involving fragmented packets.
- To avoid this security issue, configure R4 to permit only non-fragmented and initial fragmented IP packets to go out its connection to VLAN 4.

2 Points

7.2 Exploit Protection

- The network administrator has reported that several internal Windows web servers are open to a recently reported vulnerability. This vulnerability relates to a buffer overflow exploit that involves someone attempting to retrieve a URL containing ‘root.exe’.
- Until there is a patch available for the vulnerability, configure R4 to filter off all HTTP GET requests that contain ‘root.exe’ in them coming from BB2.

3 Points