

IEWB-RS-VOL2 Lab 1

Difficulty Rating (10 highest): 5

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 or add any IP addresses from the initial configuration unless otherwise specified
- 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 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	12
Frame Relay	6
HDLC/PPP	0
Interior Gateway Routing	24
Exterior Gateway Routing	9
IP Multicast	9
IPv6	9
QoS	6
Security	5
System Management	11
IP Services	9

GOOD LUCK!

Troubleshooting:

There are three issues with the initial configurations applied to the devices that will need to be resolved before the lab can be completed.

1. Bridging & Switching

1.1. VLAN Assignments

- Configure the VTP domain CISCO-A on SW1 and SW2.
- Configure the VTP domain CISCO-B on SW3 and SW4.
- SW3 should be a VTP server and SW4 should be a VTP client.
- Create and configure the VLAN assignments as follows:

Catalyst Port	Interface	VLAN	Name
SW1 Fa0/1	R1 Fa0/0	Routed	N/A
SW1 Fa0/3	R3 E0/0	33	VLAN_33
SW1 Fa0/5	R5 E0/0	105	VLAN_105
SW1 Fa0/7	N/A	28	VLAN_28
SW1 Fa0/14	SW2 Fa0/14	Routed	N/A
SW2 Fa0/2	R2 Fa0/0	28	VLAN_28
SW2 Fa0/4	R4 E0/0	45	VLAN_45
SW2 Fa0/6	R6 G0/0	46	VLAN_46
SW2 Fa0/7	N/A	28	VLAN_28
SW2 Fa0/14	SW1 Fa0/14	107	VLAN_107
SW2 Fa0/21	SW4 Fa0/18	105	VLAN_105
SW2 Fa0/24	BB2	102	VLAN_102
SW2 V28	N/A	N/A	VLAN_28
SW3 Fa0/3	R3 E0/1	Routed	N/A
SW3 Fa0/5	R5 E0/1	45	VLAN_45
SW3 Fa0/21	SW4 Fa0/21	107	VLAN_107
SW3 Fa0/24	BB3	33	VLAN_33
SW4 Fa0/4	R4 E0/1	46	VLAN_46
SW4 Fa0/6	R6 G0/1	6	VLAN_6
SW4 Fa0/18	SW2 Fa0/21	Routed	N/A
SW4 Fa0/21	SW3 Fa0/21	Routed	N/A
SW4 V102	N/A	N/A	VLAN_102

2 Points

1.2. Trunking

- Configure trunking between the switches according to the table below.

Catalyst Port	Interface	Status
SW1 Fa0/13	SW2 Fa0/13	Trunk
SW1 Fa0/15	SW2 Fa0/15	Trunk
SW1 Fa0/21	SW4 Fa0/15	Trunk
SW2 Fa0/16	SW3 Fa0/16	Trunk
SW2 Fa0/17	SW3 Fa0/17	Trunk
SW2 Fa0/18	SW3 Fa0/18	Trunk
SW3 Fa0/19	SW4 Fa0/19	Trunk
SW3 Fa0/20	SW4 Fa0/20	Trunk

- Ensure that other ports never become trunk ports.
- Trunk ports should never become access ports under any circumstance.
- Traffic from VLAN 46 should not be tagged with a VLAN header when it is sent over any trunk link with the exception of the trunk link between SW1 and SW4.

2 Points

1.3. Spanning-Tree Protocol

- Ports SW1 Fa0/7 and SW2 Fa0/7 connect to your corporate conference room. Recently your network administrator has been getting complaints that when users plug their laptops into the conference room it either takes a very long time to get an IP address from the DHCP server, or the DHCP request times out. After further investigation, you have discovered that spanning-tree convergence time is to blame.
- In order to resolve this configure SW1 and SW2 so that users in VLAN 28 do not have to wait for spanning-tree's forwarding delay when they connect to the network.

2 Points

1.4. Spanning-Tree Protocol

- After implementing the change in spanning-tree configuration for VLAN 28, one of your users plugged a switch into the conference room and crashed your entire network. After further investigation, you have discovered that a spanning-tree loop was to blame.
- In order to prevent this problem in the future ensure that any ports in VLAN 28 will be shut down if a device running spanning-tree protocol is detected.

2 Points

1.5. Traffic Engineering

- Configure the network in such a way to ensure that VLAN 102's traffic never traverses SW3.
- Additionally ensure that no other VLAN traffic follows the path that VLAN 102 does through the switched network.

2 Points

1.6. VLAN Security

- The network administrator has requested the ports SW1 Fa0/7 and SW2 Fa0/7 should not be able to communicate directly with each other within VLAN 28.
- These ports should still be allowed to communicate with R2's F0/0 interface but not SW2's V28 interface.
- You are allowed to additionally create and use VLAN 281 for this task.
- Do not use a VLAN ACL to accomplish this.

2 Points

2. Frame Relay

2.1. Hub-and-Spoke

- Using only physical interfaces configure a Frame Relay hub-and-spoke network between R1, R2, and R3 with R2 as the hub.
- Traffic from R1 destined for R3 should transit R2 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

2.2. Hub-and-Spoke

- Using only physical interfaces configure a Frame Relay hub-and-spoke network between R3, R4, and R5 with R5 as the hub.
- Traffic from R3 destined for R4 should transit R5, 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

2.3. Point-to-Point

- Using only the physical interface configure a Frame Relay point-to-point connection between R6 and BB1.
- Use static layer 3 to layer 2 resolution to reach BB1's address 54.X.1.254.

2 Points

3. Interior Gateway Routing

3.1. OSPF

- Configure OSPF area 0 on the Frame Relay connection between R3, R4, and R5.
- Ensure that R5 is always elected the Designated Router for this segment.
- Do not use the **neighbor** statement under the OSPF process to accomplish this.
- Advertise the Loopback 0 interfaces of R3, R4, & R5 into OSPF area 0.

3 Points

3.2. OSPF

- Configure OSPF area 45 on VLAN 45 between R4 and R5.
- Ensure that host devices running OSPF on this segment cannot intercept the OSPF communication between R4 and R5.

3 Points

3.3. OSPF

- Configure OSPF area 46 on VLAN 46 between R4 and R6.
- Advertise R6's Loopback 0 interface into OSPF area 46.
- Advertise VLAN 6 into OSPF on R6; do not use the **network** statement under the OSPF process to accomplish this.

2 Points

3.4. OSPF

- The Ethernet link between R4 and R5 will be used primarily as a backup of the Frame Relay circuit between them.
- Configure the network so that traffic is only sent over this Ethernet segment if the Frame Relay circuit between R4 and R5 is down.
- Do not use the **backup interface** command to accomplish this.

2 Points

3.5. OSPF

- To minimize network downtime in the event of a failure configure the network so that R4 can detect a loss of the Frame Relay circuit to R5 within 1 second.

2 Points

3.6. EIGRP

- Configure EIGRP AS 100 on R1, R2, R3, R5, SW1, SW2, and SW4.
- Configure EIGRP on the Ethernet segments as below:
 - R1 and SW1
 - R2 and SW2
 - R5 and SW4
 - SW1 and SW4
- Configure EIGRP on the Frame Relay network between R1, R2, and R3.
- Advertise the Loopback 0 interfaces of R1, R2, SW1, SW2, and SW4 into the EIGRP domain.
- Do not send EIGRP packets out any other interfaces; do not use the **passive-interface** command to accomplish this.

3 Points

3.7. EIGRP

- Advertise VLAN 33 and R3's interface E0/1 into the EIGRP domain.
- These prefixes should appear as follows throughout the EIGRP domain:

```
D EX    204.12.x.0 [170/...  
D EX    183.x.39.0 [170/...
```

2 Points

3.8. EIGRP

- Configure EIGRP AS 10 on R6.
- Enable EIGRP on the Frame Relay segment between R6 and BB1.
- Administrators of your network are concerned about false routing information being injected from the provider network. In order to ensure that all routes learned over the Frame Relay cloud are legitimate configure R6 to use the most secure authentication for any neighbor relationships formed on this interface.
- Use key number 1 with a password of CISCO for this authentication.

2 Points

3.9. RIPv2

- Configure RIP on SW4.
- Enable RIP on the Ethernet segment connecting to BB2.
- In order to protect against false route injection from RIP as well, configure SW4 to use the strongest authentication on any RIP updates received on this Ethernet segment using key 1 and the password CISCO.

2 Points

3.10. Basic Connectivity

- Configure SW3 so that all traffic for non-directly connected destinations is sent to R3.
- Devices in the network do not need reachability to SW3's Loopback 0 interface, but should have reachability to 183.X.39.9.
- Do not enable IP routing on SW3 to accomplish this.

1 Point

3.11. IGP Redistribution

- Redistribute between RIP and EIGRP on SW4.
- Redistribute between OSPF and EIGRP on R3, R5, and R6.
- R5 should route through R3 to get to R1's Loopback 0 interface.
- R5 should still be able to reach this prefix if the Frame Relay circuit between R2 and R3 is down.

2 Points

4. Exterior Gateway Routing

4.1. BGP Peering

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

Device	BGP AS
R1	200
R2	200
R3	100
R4	100
R5	100
R6	100
SW1	200
SW4	200
BB1	54
BB2	254
BB3	54

- Configure the BGP peering sessions as follows:

Device 1	Device 2
R6	BB1
R5	R3
R5	R4
R5	R6
R5	SW4
SW4	BB2
SW4	SW1
SW1	R1
R1	R2
R3	R2
R3	BB3

- The BGP peering sessions between R4 & R5 and R5 & R6 should remain up if R4 loses its connection to the Frame Relay cloud.
- Administrators of your network are concerned about insecure BGP updates being passed over VLAN 102. To resolve this configure SW4 to authenticate its BGP peering session with BB2 using the password CISCO.

3 Points

4.2. BGP Bestpath Selection

- For the purposes of load-sharing and redundancy, AS 100 has multiple connections to AS 54. In order to maximize throughput your corporate policy dictates that all traffic destined for prefixes originated in AS 54 should traverse the Frame Relay link between R6 and BB1.
- In the case that the Frame Relay link between R6 and BB1 goes down AS 100 should still have reachability to AS 54 via the Ethernet segment between R3 and BB3.
- Do not modify weight to accomplish this.

3 Points

4.3. BGP Bestpath Selection

- Configure a new Loopback interface on R1 with the IP address 150.X.11.1/24 and advertise it into BGP.
- Configure AS 200 so that all traffic from AS 100 destined to this prefix traverses the Ethernet segment between SW4 and R5.
- In the case that the link between SW4 and R5 is down traffic destined for the 150.X.11.0/24 prefix should transit the Frame Relay link between R2 and R3.
- Do not use AS-Path prepending to accomplish this.

3 Points

5. IP Multicast

5.1. PIM

- Configure IP Multicast routing on R2, R3, and R5.
- Configure PIM on the following interfaces:

Device	Interface
R2	Fa0/0
R2	S0/0
R3	E0/0
R3	S1/0
R3	S1/1
R5	E0/0
R5	S0/0

2 Points

5.2. RP Assignment

- Configure R3 to announce its most reliable interface as the RP for all multicast groups.
- R2 should be responsible for group to RP mappings.

2 Points

5.3. Multicast Testing

- There is a Windows® Media Server located on VLAN 28 that is streaming a video feed into your network, however your administrators have been getting complaints from users on VLAN 105 that they are unable to receive this feed.
- In order to help track down the source of this problem configure R5's Ethernet interface attached to VLAN 105 to join the multicast group 226.26.26.26.
- Ensure that R5 responds to ICMP echo-requests sourced from VLAN 28 which are sent to 226.26.26.26.

3 Points

5.4. Multicast Filtering

- Development engineers are testing a new multicast application located on VLAN 28 prior to its deployment in your network. This application is generating random multicast streams destined for addresses in the administratively scoped multicast range.
- In order to prevent this test traffic from being unnecessarily forwarded throughout the network configure R3 so that hosts in VLAN 33 are not allowed to join any groups in this range.

2 Points

6. IPv6

6.1. IPv6 Addressing

- The network administrator has requested that VLAN 46 and VLAN 105 be configured to support a test deployment of IPv6.
- Address R4's interface attached to VLAN 46 with the IPv6 network 2001:CC1E:X:404::/64.
- Address R5's interface attached to VLAN 105 with the IPv6 network 2001:CC1E:X:505::/64.
- The host addresses on these interfaces should be derived from the interface's MAC address.

3 Points

6.2. IPv6 Tunneling

- In order to connect these two isolated networks you have decided to tunnel IPv6 over your existing IPv4 infrastructure, however you want to ensure that this connection can survive a failure of the Frame Relay circuit between R4 and R5.
- To accomplish this configure a tunnel between R4 and R5 using their Loopback0 interfaces as the source.
- The tunnel should use the addresses 2001:CC1E:X:4545::Y/64.
- This tunnel should use a mode that specifies IPv6 as the passenger protocol and IPv4 as the encapsulation and transport protocol.

3 Points

6.3. RIPng

- Enable RIPng on VLAN 46, VLAN 105, and the tunnel interfaces.
- Use CISC0 as the identifier string for the RIPng process on both R4 and R5.
- R4 and R5 should be able to ping other's IPv6 enabled Ethernet interfaces using their respective hostnames.

3 Points

7. QoS

7.1. Frame Relay Traffic Shaping

- You have been noticing drops on R5's connection to the Frame Relay cloud. After further investigation, you have discovered that R5 has been overwhelming R3 and R4's connections to the Frame Relay cloud. Configure Frame Relay Traffic Shaping on R5 in order to resolve this issue.
- R5's connection to the Frame Relay cloud supports a transmission rate of 1536Kbps.
- R5 should send at an average rate of 128Kbps on DLCI 513 to R3.
- R5 should send at an average rate of 512Kbps on DLCI 504 to R4.
- In the case that the Frame Relay cloud notifies R5 of congestion it should reduce its sending rate to no lower than 96Kbps for the DLCI to R3 and 384Kbps for the DLCI to R4.
- In the case that R5 has accumulated credit it should be allowed to burst up to the maximum transmission rate supported on the circuit to R4.
- Bursting on the circuit to R3 should not be allowed.
- Assume an interval (Tc) of 50ms.

3 Points

7.2. Rate Limiting

- One of your NOC engineers has noticed suspiciously high utilization on the Ethernet segment of R1. After further investigation you have found that a large number of ICMP packets have been traversing this link.
- In order to alleviate congestion configure R1 so that it does not send more than 128Kbps of ICMP traffic out this interface.
- Allow for a burst of 1/4th of this rate.

3 Points

8. Security

8.1. Denial of Service Tracking

- Your network administrators have been getting complaints from users that the web server with the IP address 183.X.28.100 is inaccessible. After further investigation you have determined that this server is undergoing a TCP SYN attack.
- In order to assist in tracking down the source of this attack configure R3 and SW4 to generate a log message when HTTP SYN packets are received on VLANs 33 or 102 respectively that are destined for 183.X.28.100.
- These log messages should include the MAC address of the device which forwarded the packet onto the segment.

3 Points

8.2. Spoof Prevention

- After reviewing your log files you have determined that the DoS attack on your web server came from hosts with spoofed source addresses.
- To help prevent this type of attack in the future configure your network so that traffic will not be accepted from BB1, BB2, or BB3 if it sourced from your address space 183.X.0.0/16.

2 Points

9. System Management

9.1. RMON

- In order to help detect possible flood attacks in the future configure R2 to generate an SNMP trap when the interface input unicast packets (`ifEntry.11.1`) value rises more than 15000 per minute, and when the value falls back below 5000 per minute.
- The sampling interval should be every sixty seconds.
- When the 15000 threshold is breached an event should be generated that reads "Above 15000 for ifInUcastPkts".
- When the value falls back to 5000 an event should be generated that reads "Below 5000 for ifInUcastPkts".
- The server to send these SNMP traps to is 183.X.17.100.
- This server will be expecting the community string to be IETRAP.

3 Points

9.2. Syslog

- In order to keep track of important device notifications your corporate policy now requires that all devices send their log messages to the syslog server located at 183.X.17.100.
- Configure all devices in the network to conform to this policy.
- R1 through R6 should send log messages using facility local5.
- SW1 and SW4 should send log messages using facility local6.
- In order to ease in identifying where specific log messages are originated from ensure that all devices source their logging messages from their respective Loopback0 interfaces.

3 Points

9.3. NTP

- After implementing syslog logging your NOC engineers have noticed inconsistent timestamps on your device logs. In order to resolve this problem you have decided to maintain consistent time by implementing Network Time Protocol.
- Configure R3 and R6 to get network time from BB3 and BB1 respectively.
- Configure R1, R2, and SW1 to get network time from R3.
- Configure R4, R5, and SW4 to get network time from R6.
- R3 should fail over and get network time from R6 in the event that BB3 becomes unavailable.
- R6 should fail over and get network time from R3 in the event that BB1 becomes unavailable.

3 Points

9.4. NTP Authentication

- In order to assure that BB1 and BB3 are legitimate time sources configure R3 and R6 to authenticate them with the MD5 password CISCO.

2 Points

10. IP Services

10.1. Traffic Accounting

- Your design team would like to implement a new QoS policy using IP precedence on the Frame Relay circuit between R2 and R3. However, prior to implementing this new policy they need to know if packets transiting this link already have an IP precedence value set.
- To accomplish this configure R2 and R3 to collect usage statistics on packets with an IP precedence value and store them locally.
- R2 and R3 should store up to 50000 of these entries in their memory.

3 Points

10.2. Gateway Redundancy

- Your administrators are concerned about default gateway redundancy for the hosts located on VLAN 105. In order to allow them to survive a network failure you have assigned the virtual IP address 183.X.105.254 as the default gateway for these hosts.
- As long as R5's Frame Relay connection is up it should respond to ARP requests sent to this IP address.
- In the event that R5's Frame Relay connection is lost hosts should use SW4 as their default gateway.
- Do not use VRRP to accomplish this.
- Configure your network to reflect this policy.

3 Points

10.3. Network Address Translation

- Your operations team does not want BB3 and its customers to have specific reachability information about your network. Instead, BB3 should only have reachability to your hosts if a connection is initiated from inside your network.
- Configure R3 to reflect this policy.
- Ensure that all devices in the 183.X.0.0/16 network can successfully ping BB3.

3 Points