

Copyright Information

Copyright © 2003 - 2010 Internetwork Expert, Inc. All rights reserved.

The following publication, *CCIE Routing & Switching Lab Workbook Volume II*, was developed by Internetwork Expert, Inc. All rights reserved. No part of this publication may be reproduced or distributed in any form or by any means without the prior written permission of Internetwork Expert, Inc.

Cisco®, Cisco® Systems, CCIE, and Cisco Certified Internetwork Expert, are registered trademarks of Cisco® Systems, Inc. and/or its affiliates in the U.S. and certain countries.

All other products and company names are the trademarks, registered trademarks, and service marks of the respective owners. Throughout this manual, Internetwork Expert, Inc. has used its best efforts to distinguish proprietary trademarks from descriptive names by following the capitalization styles used by the manufacturer.

Disclaimer

The following publication, *CCIE Routing & Switching Lab Workbook Volume II*, is designed to assist candidates in the preparation for Cisco Systems' CCIE Routing & Switching Lab exam, specifically the **Configuration** portion of the Lab exam. While every effort has been made to ensure that all material is as complete and accurate as possible, the enclosed material is presented on an "as is" basis. Neither the authors nor Internetwork Expert, Inc. assume any liability or responsibility to any person or entity with respect to loss or damages incurred from the information contained in this workbook.

This workbook was developed by Internetwork Expert, Inc. Any similarities between material presented in this workbook and actual CCIE™ lab material is completely coincidental.

Table of Contents

About IEWB-RS-VOL2	2
IEWB-RS-VOL2 Lab 1	4
IEWB-RS-VOL2 Lab 2	18
IEWB-RS-VOL2 Lab 3	32
IEWB-RS-VOL2 Lab 4	46
IEWB-RS-VOL2 Lab 5	60
IEWB-RS-VOL2 Lab 6	74
IEWB-RS-VOL2 Lab 7	88
IEWB-RS-VOL2 Lab 8	106
IEWB-RS-VOL2 Lab 9	120
IEWB-RS-VOL2 Lab 10	136
IEWB-RS-VOL2 Lab 11	150
IEWB-RS-VOL2 Lab 12	166
IEWB-RS-VOL2 Lab 13	182
IEWB-RS-VOL2 Lab 14	198
IEWB-RS-VOL2 Lab 15	214
IEWB-RS-VOL2 Lab 16	228
IEWB-RS-VOL2 Lab 17	242
IEWB-RS-VOL2 Lab 18	256
IEWB-RS-VOL2 Lab 19	268
IEWB-RS-VOL2 Lab 20	282

About IEWB-RS-VOL2

Internetnetwork Expert's CCIE Routing & Switching Lab Workbook Volume II (IEWB-RS-VOL2) is designed to be used as a supplement to other self-paced and instructor-led training materials in preparation for Cisco Systems' CCIE Routing & Switching Lab Exam. Please, refer to the company website <http://www.INE.com> for more information on additional products and product bundles.

IEWB-RS-VOL2 consists of various lab scenarios, designed from the ground up, based on Cisco Systems' newest specification for the CCIE Routing & Switching Lab Exam. The labs contained in IEWB-RS-VOL2 are designed to simulate the actual CCIE Routing & Switching Lab Exam and at the same time illustrate the principles behind the technologies that it covers.

IEWB-RS Volume 2 Lab 1

Difficulty Rating (10 highest): 6

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	4
IPv4	12
IPv6	7
MPLS VPN	6
Multicast	7
Security	10
Network Services	21
QoS	12

GOOD LUCK!

1. Layer 2 Technologies

Some Layer 2 settings have been pre-configured for you; VLANs have been created and assigned to the switchports according to the diagram supplied with the task.

1.1 Layer 2 Features

- Ensure SW2 is the primary STP root bridge for VLAN105
- Configure the network in such a way to ensure that VLAN 102's traffic never traverses SW3.
- Ports Fa0/7 on SW1 and Fa0/7 on SW2 belong to VLAN 28.
- 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.
- Ensure that any ports in VLAN 28 will be shut down if a device running spanning-tree protocol is detected.
- Additionally, the same ports should not be able to communicate directly with each other within VLAN 28.
- These ports should still be allowed to communicate with R2's VLAN28 interface but not SW2's V28 interface.
- You are not allowed to use VLAN ACLs to accomplish this, but you are allowed to create and use additional VLAN 281.

4 Points

2. IPv4

Some IGP protocol settings and IP addressing have been preconfigured for you. Notice that there might be some issues deliberately introduced into the initial configurations. Use the diagram as you reference to fix those.

2.1 OSPF

- Ensure that R5 is always elected the Designated Router for the segment between R3, R4 and R5.
- Ensure that other devices running OSPF on the segment between R4 and R5 cannot intercept the OSPF communication between R4 and R5.
- Advertise VLAN 6 into OSPF on R6; but do not use the `network` or `ip ospf` statements to accomplish this.
- Configure the network so that traffic is only sent over the VLAN45 link if the Frame Relay circuit between R4 and R5 is down. Do not use the `backup interface` command to accomplish this.
- To minimize 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

4 Points

2.2 IGP Features

- Advertise VLAN 33 and R3's interface connected to SW3 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/...
```

- In order to ensure that all routes learned over the Frame Relay cloud via EIGRP 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.
- 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.
- Redistribute between RIP and EIGRP on SW4.
- Redistribute between OSPF and EIGRP on R3, R5, and R6.

3 Points

2.3 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 the above requirements.
- 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.

4 Points

3. IPv6

3.1 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.
- 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.
- Enable EIGRPv6 on VLAN 46, VLAN 105 and the tunnel interfaces.
- Use 45 as the AS# for the processes on both R4 and R5.

4 Points

3.2 IPv6 Multicast Basics

- Configure R6 to join the multicast group FF06::6 on its connection to R4.
- Configure R4 to accept MLD messages only for the group range FF06::/16.
- Use R5 as the PIM RP and ensure multicast packets from R5 can reach R6.
- Do not configure the RP address statically and do not configure the same router as a BSR and RP.

3 Points

4. MPLS VPN

4.1 LDP

- Configure MPLS label redistribution between R4, R5 and R6 using the industry-standard protocol.
- Make sure the labels are redistributed even if the primary Frame-Relay link between R4 and R5 fails.
- Configure so that the labels are only generated for the respective router's Loopback0 interfaces.

3 Points

4.2 VPN

- Using the RD 100:5 and 100:6 configure two VRFs named VPN_A and VPN_B in R5 and R6 respectively.
- Use the same route-target values to tag the respective routes.
- Create two new Loopback interfaces in R5 and R6 with the IP addresses 172.16.5.5/24 and 192.168.6.6/24 and assign them to VRFs VPN_A and VPN_B respectively.
- Configure R5 and R6 to provide reachability between the two subnets.

3 Points

5. IP Multicast

5.1 RP Assignment

- Discover the active multicast topology using the respective show commands.
- Configure R3 to announce its most reliable interface as the RP for all multicast groups using Auto-RP protocol.
- R2 should be responsible for group to RP mappings.

2 Points

5.2 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.
- You are allowed to use one static multicast route to accomplish this.

3 Points

5.3 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. Security

6.1 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

6.2 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

6.3 Information Leaking

- Your security manager is concerned with potential network reconnaissance attacks originating from behind BB2.
- In order to minimize information exposure, configure SW4 not to notify hosts behind BB2 of any networks that it does not know about.
- Additionally, SW4 should not disclose its network mask to any host on the VLAN 102 segment.

2 Points

6.4 Control Plane Protection

- Configure R4 to drop any outgoing IP packets with the TTL lower than 3.
- Log the drop events to the system console and the router's memory buffer.

3 Points

7. Network Services

7.1 RMON

- In order to help detect possible flood attacks in the future configure R2 to generate an SNMP trap when the MIB value ifEntry.11.1 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

7.2 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.
- R1 and R2 should cooperatively coordinate their clocks and adjust each other's time.

2 Points

7.3 NTP Authentication

- In order to ensure that your internal time servers are not being spoofed, configure R3 and R6 to be authenticated using the MD5 password CISCO.
- Make sure the NTP peering session between R1 and R2 is authenticated as well.

3 Points

7.4 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

7.5 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 physical connection is up it should respond to ARP requests sent to this IP address.
- In the event that R5's Frame Relay port goes down hosts should use SW4 as their default gateway.
- Do not use VRRP or GLBP to accomplish this.
- Configure your network to reflect this policy.

3 Points

7.6 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

7.7 Embedded Event Management

- Configure R6 to generate a syslog message once its Frame-Relay output utilization exceeds 80% of the total interface bandwidth.
- The sylog message should contain the interface name and the current interface load.
- Ensure your configuration responds to the transmit load changes as fast as possible with Cisco IOS routers.
- Do not use the RMON feature to accomplish this.

4 Points

8. QoS

8.1 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

8.2 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

8.3 CBWFQ

- Your company plans to reduce expenses by sending PSTN calls to the remote office connected to R5 across the WAN. Currently the WAN link is used primarily for data transfers and remote desktop application.
- Configure R5 to allocate 64Kbps of PVC bandwidth to VoIP bearer traffic, which is marked as DSCP EF.
- At the same time, guarantee 30% of remaining bandwidth to Citrix application traffic.
- Set the queue depth for the Citrix traffic class to 16 packets.
- All other remaining traffic should receive flow-based fair scheduling.

3 Points

8.4 Catalyst QoS

- R6 should see OSPF packets sent from R4 marked with IP Precedence value of 3.
- Limit the aggregate traffic rate for packets sent to R6 VLAN 46 interface to 4 Mbps but ensure the exceeding packets are queued.
- At the same time, traffic flows generated by HTTP server responses should be limited to 1Mbps on this connection.
- Use DSCP value of CS2 to mark the HTTP packets for this task.
- Do not configure any of the routers to accomplish this.

3 Points

IEWB-RS-VOL2 Lab 2

Difficulty Rating (10 highest): 6

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	8
IPv4	21
IPv6	3
MPLS VPN	4
Multicast	4
Security	12
Network Services	15
QoS	12

GOOD LUCK!

1. Layer 2 Technologies

Some Layer 2 settings have been pre-configured for you; VLANs have been created and assigned to the switchports according to the diagram supplied with the task.

1.1 Link Aggregation

- Configure interfaces Fa0/19 – Fa0/21 on SW1 and interfaces Fa0/13 – Fa0/15 on SW4 to be bound as one logical Layer 2 link.
- SW1 should actively negotiate this link through LACP.
- SW4 should respond to SW1's LACP requests, but should not initiate negotiation.
- Ensure that SW1 is the 801.2ad *decision maker* for this logical link.
- Both switches should actively negotiate ISL as the trunking encapsulation for this logical link.

3 Points

1.2 802.1x Authentication

- Your network administrator has voiced some concerns relating to the security of SW1's ports Fa0/9 and Fa0/10 which are being used in the company's public meeting room. In order to provide added security for these connections, a new corporate policy mandates that clients accessing the network through these ports must authenticate prior to being granted access to the network.
- A RADIUS server with the IP address 204.12.X.100 and the key CISCO has been configured for SW1 to authenticate these clients.
- The radius server is expecting the source of these packets to come from 150.X.7.7.
- Do not use authentication on any of terminal lines of SW1.

3 Points

1.3 Performance Optimization

- Your company has recently acquired another company and will be merging your network with the other network in the near future.
- The other company's network currently contains 4,000 OSPF routes.
- To ensure problems do not arise when the two networks are merged, configure SW1 and SW2 so that their routing tables can support the 4,000 routes from the new company's network plus the existing unicast routes that will exist in your network.

2 Points

2. IPv4

Some IGP settings along with IP addressing have been preconfigured for you in this scenario.

2.1 OSPF

- Ensure there is no DR or BDR election on the segment between R1, R2, R3 and R4.
- Do not use the `neighbor` statement to accomplish this.
- Authenticate the adjacency between R1 and SW1 using the clear-text password CISCO.

3 Points

2.2 EIGRP

- Enable EIGRP on VLAN 26 between R2 and R6.
- As an added security measure, configure the network so that hosts running EIGRP on VLAN 26 cannot intercept the EIGRP communication between R2 and R6.
- Configure R6 so that the rest of the routers in the EIGRP domain have only one route to the prefixes learned from BB1 with a first octet of 200.
- This route should not unnecessarily overlap IP address space.

3 Points

2.3 RIP Filtering

- Configure SW1 so that it does not accept routes with an even second octet from BB3.
- The access-list used to accomplish this should not contain more than one line.
- Do not use the `distribute-list` or `distance` keywords to accomplish this.

2 Points

2.4 IGP Redistribution

- Redistribute between RIP and OSPF on SW1.
- Redistribute between OSPF and EIGRP on R2, R3, and R4.
- Ensure that full reachability is maintained throughout the IGP domain when the Frame Relay circuit between R3 and R5 is down.

3 Points

2.5 BGP Peering

- Configure a BGP peering session between R5 and BB2.
- As you are concerned about false routing information being injected from unauthorized sources configure R5 to authenticate its BGP peering session with BB2 using the password CISCO.
- AS 400 has recently acquired from AS 100, but due to AS 54's change control policy BB3's configuration will not be modified until AS 54's normally scheduled maintenance window. As an interim solution, configure AS 400 so that BB3 still thinks SW1 is still in AS 100.

3 Points

2.6 BGP Filtering

- Part of the acquisition agreement between AS 400 and AS 100 stipulates that AS 400 will not provide transit for traffic coming from AS 54 and its customers that is destined for AS 254. Configure AS 400 to reflect this policy.

2 Points

2.7 BGP Summarization

- In order to facilitate in keeping the global BGP table as small as possible your BGP routing policy dictates that R5 should advertise one route representing your entire major network 132.X.0.0/16 to BB2.
- Furthermore, since the Ethernet segment between R5 and BB2 is AS 254's only connection to your network, AS 254 does not need to have any longer matches than the summary route.
- Configure your network to reflect this policy, but do not allow any other routers within your network to see the summary route.
- Do not apply any access-list or prefix-list filtering towards AS 254 to accomplish this.

3 Points

2.8 BGP Tuning

- In order to increase BGP routing stability, configure R5 to respond to IGP events with 15 seconds delay.
- At the same time, to improve BGP convergence, configure R5 to batch and send routing updates to BB2 every 3 seconds.

2 Points

3. IPv6

3.1 IPv6 Deployment

- A new corporate directive has mandated that IPv6 be deployed company-wide within the next six months. However your network administrator is concerned that when IPv6 is deployed over the WAN there may be problems with running it over the Frame Relay full-mesh network between R1, R2, R3, and R4. Therefore the administrator has requested for you to configure IPv6 on the Frame Relay connection between R2 and R3 to determine if there will be any issues with the deployment.
- Configure R2 and R3 with the addresses 2001:CC1E:X::Y/128 under their respective Loopback 0 interfaces.
- Use 2001:CC1E:X:2323::Y/64 for the Frame Relay network.
- A static route pointing to each other's /128 Loopback addresses is permitted on R2 and R3 to test reachability.

3 Points

4. MPLS VPN

4.1 L2 VPN

- Configure R4 and R6 to provide transparent Layer-2 VPN services between VLAN6 and VLAN4 attached circuits respectively.
- Use AToM as the encapsulation method to accomplish this task.
- You are allowed to use two static routes and create additional interfaces to accomplish this.

4 Points

5. Multicast

5.1 Multicast Testing

- Discover the active multicast topology using the respective show commands.
- Configure R2's most reliable interface as the rendezvous-point (RP) for all multicast groups.
- In order to facilitate in testing reachability throughout the multicast domain your configure SW1's VLAN interface participate in the multicast group 228.28.28.28.
- Ensure that SW1 responds to ICMP echo-requests sent from VLAN 26 to this multicast group address.

2 Points

5.2 Multicast Traffic Control

- During the testing phase of your multicast deployment one of your network administrators has reported that R3 has been receiving traffic source from VLAN 26 destined for the multicast group 228.28.28.28 even though it does not have any attached members.
- After further investigation you have found that R1 is also receiving traffic from R2 for feeds destined for R3. Configure R2 to resolve this problem.

2 Points

6. Security

6.1 Router Hardening

- After returning from a network security class one of your network administrators has convinced your manager that R5 is open to a variety of security vulnerabilities. To say the least, your manager is not happy that these vulnerabilities have been left unchecked for so long.
- In order to appease your manager configure R5 to conform to the following security recommendations:
 - Drop all source routed packets
 - Disable CDP and proxy ARP on the Ethernet segment to BB2
 - Disable BOOTP server
- A banner message should be displayed to all users that telnet into the router that states: "Access to this device or the attached networks is prohibited without express written permission. Violators will be shot on sight".

2 Points

6.2 Zone-Based Firewall

- Your NOC engineers have noticed that R2 and R4 are being polled via SNMP from an unauthorized source.
- To avoid any problems associated with unauthorized polling configure your network so that all SNMP requests received from behind BB1 and BB2 are filtered out by R5 and R6 respectively.
- Use Zone-Based Firewall feature to accomplish this.

4 Points

6.3 Traffic Logging

- After blocking SNMP from outside the network it appears that the device polling R2 and R4 is internal. In order to help track down the source of the SNMP packets configure R2 and R4 to generate a log message whenever a device attempts to poll them using the Read-Only community string of 'public'.
- These log messages should be sent to the syslog server at 132.X.33.100.

3 Points

6.4 ICMP Filtering

- Recently your network team thwarted an attempted SMURF attack issued by a disgruntled ex-administrator. In order to prevent this type of attack in the future you have decided to limit the amount of ICMP traffic that R1 permits inbound on its interface attached to SW1.
- Configure your network so that ICMP traffic is only allowed into your network via the connection between SW1 and R1 if the traffic was initiated from behind R1.
- For diagnostic and troubleshooting purposes ensure that users throughout your network are still able to traceroute from behind R1.

3 Points

7. Network Services

7.1 RMON

- After the network was overwhelmed by the latest Microsoft® worm your network administrators have requested that R5 and R6 be configured to generate an SNMP trap whenever their average five minute CPU utilization (`1system.58.0`) reaches 75%.
- The sampling interval should be done once per minute.
- When the 75% threshold is breached, an event should be generated that reads "Five Minute CPU Average Above 75%".
- When the utilization falls back below 40%, an event should be generated that reads "Five Minute CPU Average Below 40%".
- The SNMP server to send these traps to is 132.X.33.100.
- The SNMP server is expecting the community string to be IETRAP.

3 Points

7.2 Remote Access

- Your manager has requested that R4 be configured to allow the users from the company's NOC to telnet in to manage the router.
- The users will expect the username to be NOC and the password to be CISCO.
- All telnet sessions should be disconnected from the router after 5 minutes of inactivity.
- The maximum amount of time a user should be allowed to telnet into R4 before being disconnected should be 15 minutes.
- Sixty seconds prior to automatically logging this user off R4 should send the user a warning message in order to give the user time to finish up and save any changes to the configuration.

3 Points

7.3 Remote Access Security

- In order to increase the security of your password database configure R4 so that the password for the NOC username is stored as an MD5 hash that represents the password CISCO.

3 Points

7.4 Syslog

- Configure R3 to log all severity 7 and below messages to a syslog server with an IP address of 132.X.33.100.
- R3 should not generate a log message when its interface Serial1/0 changes status, but should generate a log when a Frame Relay DLCI changes status.

3 Points

7.5 System Management

- Your security manager is concerned with the possibility of unauthorized personnel gaining physical access to the networking equipment.
- In order to reduce the risk of a direct physical device exposure, ensure that no one could reload your switches and start the initial configuration using the “Mode” button on the front panel.

3 Points

8. QoS

8.1 Congestion Management

- Users behind BB2 accessing an SMTP server at 132.X.3.100 have been complaining about slow response time. After further investigation, you notice congestion on the Frame Relay link between R3 and R5.
- Configure R3 and R5 so that SMTP packets to and from the server are guaranteed at least 256Kbps during times of congestion across the Frame Relay link.
- Assume that both R3 and R5 both have a port speed of 512Kbps.

3 Points

8.2 Policy Routing

- Users in VLAN 26 have been complaining about slow response time to an FTP server located in VLAN 33 with the IP address 132.X.33.33.
- In order to alleviate congestion to and from this server a new traffic engineering policy dictates that all FTP traffic coming from VLAN 26 destined for this server and back should use the Serial link between R2 and R3 whether the link is up, irrespective of IGP routes learned.
- All other traffic destined to this server should follow normal forwarding.
- Assume that this FTP server does not FTP PASV command.
- You are allowed to use policy routing to accomplish this.

3 Points

8.3 Congestion Management

- As an additional measure to reduce FTP response time ensure that bidirectional FTP traffic between VLAN 26 and the FTP server is guaranteed a least 256Kbps on the Serial link between R2 and R3 during times of congestion.
- Assume the link physical rate to be 1536Kbps that is T1 equivalent.

3 Points

8.4 Frame Relay Traffic Shaping

- You have noticed a large number of frames arriving on R2's DLCI 204 and R4's DLCI 402 with the Frame Relay DE bit set. Since voice traffic will be transferred across these DLCIs in the near future your corporate policy now mandates that R2 and R4 must conform to the provider's CIR.
- Configure Frame Relay Traffic Shaping on R2 and R4 in order to resolve this issue using the following information:
 - R2's connection to the Frame Relay cloud has a port speed of 512Kbps
 - R2 has had a CIR of 128Kbps provisioned for DLCI 204 by the telco
 - R4's connection to the Frame Relay cloud has a port speed of 512Kbps
 - R4 has had a CIR of 128Kbps provisioned for DLCI 402 by the telco
- Under no circumstances should R2 or R4 burst above the provisioned CIR for these DLCIs.
- To ensure that you are not oversubscribing your Frame Relay circuits limit every other utilized DLCIs on R2 and R4 to a rate equal to half of the remaining bandwidth (port speed minus the provisioned CIR) of the circuit between R2 and R4.
- Use the lowest interval (Tc) available for the DLCIs between R2 and R4.
- All other DLCIs should use the default interval (Tc).

3 Points

IEWB-RS-VOL2 Lab 3

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	6
IPv4	17
IPv6	6
MPLS VPN	9
Multicast	6
Security	9
Network Services	19
QoS	7

GOOD LUCK!

1. Layer 2 Technologies

Some Layer 2 settings have been pre-configured for you; VLANs have been created and assigned to the switchports according to the diagram supplied with the task

1.1 IP Bridging

- R1 and R3 are in the same IP subnet, but in different broadcast domains.
- Configure R6 to bridge IP traffic between VLAN 16 and VLAN 36.
- Ensure that the rest of the routing domain can communicate with both R1 and R3 via IP.

3 Points

1.2 Spanning-Tree Protocol

- Configure SW1 as the spanning-tree root for VLANs: 4, 44, 52, and 63.
- All traffic between SW1 and SW2 for these VLANs should transit the trunk between SW1 and SW2's port Fa0/15.
- In the case that port Fa0/15 goes down, traffic for these VLANs should transit port Fa0/14.
- As a last resort, traffic for these VLANs should transit port Fa0/13 if both of the other trunk links are down.
- The above configuration should be done on SW1.
- In order to minimize network downtime in the event of a failure, configure SW2 so that traffic continues forwarding within three seconds if either port Fa0/15 or Fa0/14 goes down.
- This should be accomplished while running PVST.

3 Points

2. IPv4

Some IGP settings and IP addressing have been preconfigured for you.

2.1 OSPF

- Ensure that R2 uses R5 as the next hop to reach R4, and vice versa.
- Advertise the Loopback 0 interfaces of R1, R2, R4, and R5 into OSPF area 0. These routes should appear with a subnet mask of /24 with the exception to the prefixes that you need to appear as /32.
- Configure OSPF area 45 on the Serial link between R4 and R5.
- This link will be used primarily as a backup of the Frame Relay circuit between R4 and R5. Configure the network so that reachability is maintained over the Serial link when R4's connection to the Frame Relay cloud is down.
- Traffic should not be routed across the Serial link when the Frame Relay circuit from R4 to R5 is up.
- Do not use the `backup interface` command to accomplish this.
- You are concerned about false routing information being injected into OSPF area 0. In order to verify the legitimacy of routing information, configure all area 0 adjacencies to be authenticated with a secure hash value of the password CISCO.

4 Points

2.2 IGP Features

- Configure the OSPF domain to reflect the following metric calculations:

Bandwidth (Mbps)	OSPF Cost
10,000	2
10	2000
1.544	12953
0.768	26041

- R5's OSPF database is growing quickly, and the router spends considerable time on the database maintenance.
- Configure R5 so that OSPF LSA are grouped, checksummed, and max aged six times more often than by default.
- In order to protect R5 against flooding with the same LSA during network instability times, ensure it holds for twice the default interval before accepting the same LSA again from its peers.
- Redistribute where necessary to obtain full IP reachability to all advertised networks.
- R5 should route through R1 to get to the prefixes learned from BB1.
- R5 should route through R2 to get to the prefixes learned from BB3.
- Ensure R5 prefers to reach prefixes learned from BB2 via RIP and not BGP.

4 Points

2.3 BGP Path Manipulation

- Configure R6 so that AS 100 cannot be used as a transit to reach prefixes in AS 54.
- Advertise VLAN 3 into BGP on R3 and configure AS 100 so that AS 200 routes through AS 300 to get to this prefix.

4 Points

2.4 BGP Attributes

- Advertise VLAN 29 into BGP on R2.
- R5 should see this prefix as follows:

```
Network          Next Hop          Metric LocPrf Weight Path
*> 136.X.29.0/24  XXX.X.XXX.X      0          100 300 i
```

- The configuration to achieve the above should not affect any other prefixes on R5.
- Configure R2 so that all traffic destined for VLAN 29 comes in the Serial link to R3.
 - In the case that this link between is down, VLAN 29 should still be accessible via the Frame Relay link.
 - None of BGP attribute manipulations in AS 100 should affect this path selection.
- Configure SW3 to advertise the EtherChannel link into BGP.
- Ensure R3 and SW3 will accept BGP updates with AS 100 in the AS path. Do not alter R2's configuration to accomplish this.

5 Points

3. IPv6

3.1 IPv6 Addressing

- Address R2's VLAN29 interface with the network 2001:CC1E:X:202::/64
- Address R4's VLAN4 interface with the network 2001:CC1E:X:404::/64.
- The host portion of the IPv6 addresses should be based partly off of their interfaces' respective MAC addresses.
- Enable communication between VLAN 29 and VLAN 4 using an IPv4 based GRE tunnel.
- Use any site-local network for the IPv6 addressing within the GRE tunnel.

3 Points

3.2 IPv6 Routing

- Using OSPFv3 provide IPv6 routing information exchange between R2 and R4.
- R4 should learn R2's IPv6 prefixes and should be sending a route to R2 allowing R2 to reach any destination.
- Ensure the loss of an OSPFv3 neighbor is detected within 3 seconds.

3 Points

4. MPLS VPN

4.1 Label Exchange

- Configure label exchange between R4 and R5 using Cisco's legacy protocol.
- Make sure the TCP session does not use the Loopback0 interfaces as sources.
- Ensure labels are exchanged on both links connecting R4 and R5.

3 Points

4.2 MPLS VPN

- Configure VLAN 57 and VLAN 44 interfaces on R5 and R4 in the VRF VPN_AB
- Use the RD value of 100:47 and two different route-target values for every VRF.
- Make sure you can ping the directly connected interfaces across the VPN cloud.

3 Points

4.3 PE-CE Routing

- Using OSPF process numbers 44 and 77 on R4 and R5 respectively, configure PE/CE routing with the respective CE devices.
- Make sure every site sees the other site's routes as inter-area summary prefixes, not an external routes.

3 Points

5. Multicast

5.1 Multicast Forwarding

- Discover the active multicast topology using the IOS show commands.
- A client located on VLAN 29 has been configured to listen for the multicast group 228.22.22.22 for testing purposes. The application used to receive the multicast feed does not support IGMP.
- Configure the network so that this host can receive traffic sent to this group.
- Ensure R2 can fast switch traffic for this group out to VLAN 29.

2 Points

5.2 Multicast Filtering

- It has come to your attention that users in VLAN 4 have been abusing your Internet connection by streaming video and audio feeds during work hours. In order to prevent this unnecessary drain on your network resources, your manager has requested for you to only allow users in VLAN 4 to receive feeds for groups that are used for business related activities.
- These groups are 225.25.25.25 and 226.26.26.26.
- Configure your network to reflect this policy.

2 Points

5.3 Multicast Filtering

- Recently, you have noticed suboptimal forwarding of multicast feeds throughout your network due to problems in your unicast routing. In order to prevent multicast feeds from looping around the network, configure R1 so that it does not send any multicast traffic out its FastEthernet interface that has a TTL of less than 13.

2 Points

6. Security

6.1 Traffic Filtering

- The network administrator has requested that R6's connection to BB1 be secured to prevent unauthorized access into your network.
- Configure R6 so that it only allows TCP, UDP and ICMP traffic in from BB1 if it was originated from behind R6, or is required for another section of the lab to work.
- Ensure that users behind R6 can still `tracert` to hosts beyond the Frame Relay cloud.

3 Points

6.2 DoS Prevention

- Users are complaining about slow response time to a web server at IP address 136.X.4.100. After further investigation, it appears that the web server is undergoing a TCP SYN flood.
- In order to help deal with these attacks, configure R4 to send a TCP reset to the web server for any TCP sessions that fail to reach the established state after 15 seconds.

3 Points

6.3 DHCP Security

- Configure R6 as a DHCP server for VLANs 16 and 36 but allocate only the IP addresses in range 136.X.136.100-136.X.136.200.
- The default gateway allocated to the hosts should be the GLBP group's virtual IP address.
- Configure SW1 and SW3 to drop unauthorized DHCP packets received on VLAN16 and VLAN36 ports.
- At the same time, R6 should populate its ARP table based on DHCP database contents.
- Rate-limit DHCP traffic on the ports connecting to R1 and R3 to 10 packets per second.

3 Points

7. Network Services

7.1 IOS Management

- Since some of your network administrators do not understand how to use the IOS CLI, they have requested that R4 be setup to be managed via HTTP. In order to minimize the risk of managing R4 through HTTP, use the following parameters:
 - Use TCP port 8080
 - Only permit access from the 136.X.2.0/24 subnet
 - Authenticate users using local username WEB and the password CISCO
 - This password should be stored in the router's configuration as an MD5 hash.

2 Points

7.2 File Management

- The NOC has reported that R1 has been having problems with its flash memory, and has been trying to load the default IOS image named "cisco2-C2600" via TFTP. In response to this the NOC has loaded the image "c2600-iuo-mz.122-13.bin" into R3's flash in case of a failure of R1.
- Configure the network so that R1 can boot this image from R3 if its flash fails again.

2 Points

7.3 Auto-Install

- A new router will be installed on the Frame Relay cloud connecting to R5 shortly using DLCI 555. This new router will need to get its configuration from a TFTP server located in VLAN 29.
- Configure R5 to use the 136.X.5.0/30 subnet for communication with the new router and provide it with IP address 136.X.5.2 via BOOTP.

3 Points

7.4 Local Authorization

- Following a recommendation by an outside consultant, management has requested that R2's default privilege level for telnet access be set to 0.
- The only commands (other than privilege 0 commands) that these users should be allowed to issue are ping and traceroute.
- If the users need privilege level 1 commands, they should be required to authenticate with the password CISCO prior to being given access.

3 Points

7.5 Local Authorization

- The first level support engineers from the company's NOC have complained to management that they are unable to troubleshoot RIP issues because they do not have enable access to R5. In response to this, management has decided that the NOC users should be able to turn on and disable RIP debugging, but not be allowed any other access.
- The NOC users will be entering R5 in user mode (privilege level 1).

3 Points

7.6 Switch Management

- Configure SW1 and SW2 to be managed via SNMP using the following parameters:
 - Contact: CCIE Lab SW1.
 - Location: San Jose, CA US.
 - Chassis ID: 221-787878.
- The network management station's IP address is 136.X.2.100, and it will be expecting the RO community string to be CISCORO and the RW community string to be CISCORW.
- SW1 and SW2 should generate SNMP traps for changes related to VTP, using the community string CISCOTRAP.

2 Points

7.7 GLBP

- Configure R1, R3 and R6 as a single virtual router using the IP address of 136.X.136.254.
- Authenticate all packet exchange using secure authentication scheme and a password of CISCO.
- Hosts on VLAN16 and VLAN36 using the virtual gateway should load-balance between R1, R3 and R6 using the proportion 1:3:6.
- R6 should be responsible for ARP responses and R1 should take its role in case of failure.
- Make sure R6 is not used for packet forwarding once it loses connectivity to AS54 via BB1.
- Use ICMP echo packets to track reachability. Send ICMP echo probes every second.

4 Points

8. QoS

8.1 Frame Relay Traffic Shaping

- The network administrator has request that Frame Relay Traffic Shaping be configured on R1, R2, R4, and R5, according to the following requirements:
 - Data should be sent at a sustained rate of 256Kbps per DLCI.
 - In the event of congestion notification, fallback to no lower than 192Kbps.
 - The routers should interpret FECN signals and back-pressure the sender by using the BECN bits.

2 Points

8.2 Rate Limiting

- In order to ensure that users on VLAN 44 are being productive during work hours, your management has requested that all HTTP responses sent out R4's interface Fa0/1 be limited to 256Kbps between the hours of 8am to 5pm Monday through Friday.
- Configure R4 to reflect this policy.

2 Points

8.3 Signaling

- Recently, you have been receiving complaints from users on VLANs 4 and 52 about low VoIP quality across the data network. After further investigation, you have determined that too much of the Frame Relay circuit between R4 and R5 is being consumed by data traffic.
- In attempt to improve VoIP performance, your network administrators have configured the client applications on these VLANs to request bandwidth reservations of the network in the transit path.
- Configure R4 and R5 to support this new setup.
- Assume that each call can reserve up to 64Kbps, and that no more than 128Kbps can be reserved at any given time.

3 Points

IEWB-RS-VOL2 Lab 4

Difficulty Rating (10 highest): 7

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	18
IPv4	6
IPv6	6
MPLS VPN	10
Multicast	7
Security	11
Network Services	14
QoS	9

GOOD LUCK!

1. Layer 2 Technologies

Some Layer 2 settings have been pre-configured for you; VLANs have been created and assigned to the switchports according to the diagram supplied with the task

1.1 Traffic Control

- Enable pruning globally within the VTP domain.
- Although SW1 and SW3 do not have VLAN 8 locally assigned, ensure that they receive unknown unicast, broadcast, or multicast traffic for VLAN 8 over their lowest numbered trunk link to SW2.
- Traffic for VLAN 8 should not be received over any of the other trunk links.

2 Points

1.2 Spanning-Tree Protocol

- Configure SW1 as the spanning-tree root for VLAN 258 and configure SW3 to become the spanning-tree root for VLAN 258 in the event SW1 is not longer available.
- All VLAN 258 traffic from SW2 to SW1 should transit SW4.
- In the event that SW2's path to SW1 through SW3 is down, SW2 should use the directly connected trunk links to reach SW1 directly.
- Use the fewest number of commands to accomplish this task and do not alter SW1's port-priorities.

3 Points

1.3 Link Failure Detection

- Administrators of your network are concerned about SW1 and SW2 not being able to detect bidirectional link failure on port Fa0/15.
- Configure SW1 and SW2 so that port Fa0/15 is brought down in the case that either switch can send traffic, but not receive, or vice versa.
- As an additional precaution, configure SW1 so that interface Fa0/15 is not mistakenly elected as a designated port in the above case.

3 Points

1.4 Spanning-Tree Protocol

- Ensure the output of the `show spanning-tree vlan 258` command on SW3 matches the highlighted output below:

```
Rack1SW3#show spanning-tree vlan 258
```

```
VLAN0258
Spanning tree enabled protocol ieee
```

```
<output omitted>
```

Interface	Role	Sts	Cost	Prio.	Nbr	Type
Fa0/5	Desg	FWD	100	128.5		Shr
Fa0/16	Desg	FWD	19	128.16		P2p
Fa0/17	Desg	FWD	19	128.17		P2p
Fa0/18	Desg	FWD	19	128.18		P2p
Fa0/19	Altn	BLK	19	128.19		P2p
Fa0/20	Altn	BLK	19	128.20		P2p
Fa0/21	Root	FWD	19	128.21		P2p

- Do not make any changes to SW3 to accomplish this task.

2 Points

1.5 Rate-Limiting

- After monitoring the inbound utilization on SW1 Fa0/1, the network administrator has requested that SW1 be configured to limit the amount of unicast traffic received from R1.
- It has been determined that the average packet size is 954 bytes and the average number of packets is 250 per second.
- Although the average packet size is 954 bytes, the network administrator has voiced concerns that the solution should accommodate various packet sizes.
- Configure SW1 to meet these requirements using the minimal amount of commands.

3 Points

1.6 QoS

- The network administrator will be implementing QoS on SW2 in the near future and would like SW2 to be configured as follows:
 - IP Precedence 0 mapped to DSCP 0
 - IP Precedence 1 mapped to DSCP 0
 - IP Precedence 2 mapped to DSCP 0
 - IP Precedence 3 mapped to DSCP 0
 - IP Precedence 4 mapped to DSCP 32
 - IP Precedence 5 mapped to DSCP 40
 - IP Precedence 6 mapped to DSCP 0
 - IP Precedence 7 mapped to DSCP 0

3 Points

1.7 QoS

- Configure SW2 to match the output of the following command:

```
Rack1SW2#show mls qos interface fa0/2
FastEthernet0/2
trust state: trust ip-precedence
trust mode: trust ip-precedence
trust enabled flag: ena
COS override: dis
default COS: 0
DSCP Mutation Map: Default DSCP Mutation Map
Trust device: none
qos mode: port-based
```

2 Points

2. IPv4

2.1 OSPF

- Administrators of your network have reported that SW2 is low on memory. After further investigation, you have determined that a large routing table is consuming the majority of SW2's memory.
- In order to cut down on the size of SW2's routing table, configure the network so that routers in OSPF area 1 do not see any inter-area or external OSPF routes.
- Ensure that devices in area 1 maintain full reachability to the rest of your network.

3 Points

2.2 OSPF

- Configure OSPF area 2 on the Ethernet, Frame Relay, and Serial segments between R4 and R5.
- Advertise the Loopback 0 interfaces of R4 and R5 into OSPF area 2.
- You are allowed to add one additional IP subnet to accomplish this.
- Mutually redistribute between RIP and OSPF on R3.

3 Points

3. IPv6

3.1 OSPFv3

- Configure IPv6 on the Frame Relay link between R1 and R2 using the addresses 2001:141:X:12::Y/64 where X and Y are your rack and router numbers in *decimal* notation.
- Advertise this link into OSPFv3 area 0.
- Configure IPv6 on the Ethernet link between R2, R5, and SW2 using the addresses 2001:141:X:25::Y/64.
- Advertise this segment into OSPFv3 area 1.

3 Points

3.2 OSPFv3 Summarization

- Create two Loopback interfaces in R5 and SW2 with IPv6 addresses 2001:150:X:Y::Y/64 and advertise them into OSPFv3 area 1
- Configure R1 to advertise a summary encompassing both Loopback subnets to R2.
- This summary should overlap the minimum amount of address space necessary.

3 Points

4. MPLS VPN

4.1 PE-CE Routing

- Customer routers BB1 and BB3 are in VPN named VPN_A and connect to R6 and R4 respectively.
- Enable RIP on the links connecting BB1 and BB3 to R6 and R4.
- Routes learned from BB1 that have an even third octet should be seen with a metric of 10 on R6.
- The access-list used to accomplish this should only have one line and should be effective for any additional networks learned from BB1 in the future.

3 Points

4.2 VPN Tunneling

- Allow for MPLS label exchange between R4 and R6.
- Use LDP and do not rely on dynamic routing for routing to the LDP endpoints.
- Ensure RIPv2 prefixes are exchanged across the VPN with metric preserved.
- You are allowed to create additional interfaces and static routes to accomplish this task.

4 Points

4.3 BGP

- Enable BGP as the PE/CE protocol between R6, R4 and BB1, BB3.
- BB3 expects R4 to be in AS 100.
- Make sure you configured PE routers so that the respective CE's will not reject the advertised BGP prefixes due to AS_PATH loops.

3 Points

5. IP Multicast

*You need to discover the active multicast configuration using the relevant **show** commands prior to starting this section. Do not change any of the existing configurations.*

5.1 Auto-RP

- Configure R2 to announce its Loopback 0 interface as a candidate rendezvous-point (RP) via Auto-RP for the multicast groups 225.0.0.0 through 225.255.255.255.
- Configure R5 to announce its Loopback 0 interface as a candidate rendezvous-point (RP) via Auto-RP for the multicast groups 239.0.0.0 through 239.255.255.255.
- SW2 should be responsible for the group to RP mappings.

3 Points

5.2 Multicast Testing

- Configure R3's VLAN37 interface as a member of the multicast group 225.25.25.25 and interface Fa0/1 as a member of 239.39.39.39.
- Ensure that R3 responds to pings sent to these multicast groups from VLANs 12 and 45.

2 Points

5.3 Multicast Rate Limiting

- In order to reduce the impact of this application on your network, configure SW1 so that no more than 1Mbps of multicast traffic is sent out towards R3.

2 Points

6. Security

6.1 Traffic Filtering

- Configure a filtering policy on R6 to conform to the following requirements:
 - Apply filtering to the VPN traffic exchanged between R4 and R6.
 - The Frame-Relay connection should be the outside interface.
 - Permit ICMP packets across the firewall (either direction).
 - Permit HTTP and SSL access to a Web server at 204.12.X.100.
 - Permit any TCP and UDP sessions initiated from behind R6 to return.
 - Limit the aggregate rate of DNS and ICMP packets inbound to 128Kbps.
- Use the Zone Based Firewall syntax to accomplish this task and apply the most secure inspection rules where possible.

6 Points

6.2 Spoof Protection

- Configure R4's VLAN43 interface to protect against spoofed IP packets originated from within the VPN address range.
- Do not use any access-lists to accomplish this task.

2 Points

6.3 Infrastructure Security

- Recent network attacks from behind BB1 prompted you to think of improving the border router security.
- In order to protect against flooding attacks targeted at R6, configure the router to limit the average rate of packets going to router's CPU to 1000 per second.
- Ensure that your configuration does not affect BGP peering sessions and router management via SSH and Telnet.

3 Points

7. Network Services

7.1 SNMP

- Configure R3 and R6 to be managed via SNMP.
- The first network management server's IP address is 141.X.7.100 and second network management server's IP address is 141.X.77.100.
- Both network management servers will be expecting the RO community string to be CISCORO, the RW community string to be CISCORW, and the community string CISCO to be used for traps.
- The first network management server will be using SNMPv1 and the second SNMP server will be using SNMPv2c.
- R3 and R6 should generate SNMP traps for changes relating to HSRP status, but these traps should only be sent to the second network management server.

3 Points

7.2 IOS Menu

- The first level support engineers from the company's NOC need to have access to R2 to ping and traceroute to R5 and R6's Loopback 0 interfaces. Since these users do not have any knowledge of Cisco IOS, the network administrator has requested that a menu be configured on R2.
- This menu should enable the NOC users to ping and traceroute to R5 and R6's Loopback 0 interfaces.
- The menu should be activated whenever the user NOC logs in using the password CISCO.
- Ensure that the NOC users can exit the menu, but do not allow them to have access to the CLI when they do so.

3 Points

7.3 DNS

- One of your network administrators has added a DNS entry that allows the NOC users to telnet to R2 by name. However, this administrator has entered the entry in DNS incorrectly to resolve to 141.X.0.22.
- Without applying this IP address to any interface permit users to telnet to R2 using this DNS entry or IP address.
- Do not use NAT to accomplish this task.

2 Points

7.4 Gateway Redundancy

- Recently, hosts on VLAN 36 suffered hours of downtime due to a hardware failure on R6. This problem was not resolved until the DHCP servers were updated to assign R3 as the default gateway for this segment. In order to prevent this problem in the future, your network team has configured half of the hosts on VLAN 36 to default to R3 and the other half to default to R6.
- Configure the network so that in the event that either R3 or R6 become unavailable, hosts on this segment should still have access to the rest of the routing domain.
- Do not use VRRP for this scenario.

3 Points

7.5 Failure Message

- Configure R3 to display a “Host Failed” message of “Connection Unsuccessful” when a telnet session to R4’s Loopback 0 interface fails.

3 Points

8. QoS

8.1 Congestion Avoidance

- Utilization monitoring on R1’s Ethernet segment has been indicating periods of high congestion followed by periods of low utilization. After further investigation you have determined that various TCP applications throughout the network are bursting and then backing off at the same time.
- In order to prevent this behavior, configure R1 to start dropping packets with an IP precedence of routine on this link when there are at least 15 packets in the output queue.

2 Points

8.2 Congestion Management

- Users in VLAN 45 have been complaining that it is taking a very long time to send e-mail messages through their SMTP server located in VLAN 258. After further investigation, you have determined that large file downloads from an FTP server are to blame.
- In order to reduce this slow response time, configure R5 so that all SMTP packets are guaranteed at least 1.5Mbps of the output queue on VLAN258 interface.
- Do not use an access-list to classify traffic for this task.

3 Points

8.3 Rate Limiting

- Configure R5 so that packets over 1250 bytes are limited to 2.5Mbps outbound on its VPN connection to BB3.

2 Points

8.4 Link Efficiency

- Configure R4 and R5 to maximize efficiency on the PPP link by guessing character streams in frames sent over the link.

2 Points

IEWB-RS-VOL2 Lab 5

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	6
IPv4	15
IPv6	8
MPLS VPN	0
Multicast	7
Security	14
Network Services	18
QoS	11

GOOD LUCK!

1. Layer 2 Technologies

1.1 EtherChannel

- Configure an EtherChannel link between SW1's interfaces Fa0/13 and Fa0/14 and SW2's interfaces Fa0/13 and Fa0/14. Use port channel number 12.
- Configure an EtherChannel link between SW1's interfaces Fa0/16 and Fa0/17 and SW3's interfaces Fa0/13 and Fa0/14. Use port channel number 13.
- Configure an EtherChannel link between SW1's interfaces Fa0/19 and Fa0/20 and SW4's interfaces Fa0/13 and Fa0/14. Use port channel number 14.
- Do not run PAgP or LACP on these links.
- All traffic sent over these trunk links should be tagged with a VLAN header.
- Do not issue any global configuration commands to accomplish this task.

3 Points

1.2 Layer 2 Services

- You have noticed very high utilization on the interface Fa0/13 between SW1 and SW2 and have determined that the majority of the traffic transiting this link is coming from a single file server located behind BB2.
- Traffic is sourced to multiple clients behind R1 with traffic patterns being unidirectional from the server to the clients.
- Configure the network in such a way that traffic sent over this EtherChannel link is distributed more evenly while taking into account the single server and multiple clients.
- Administrators of your network have noticed that some traffic has been leaking between VLAN 8 and VLAN 88. After further investigation, you have determined that SW2's CAM table is maxed out and has been treating some unicast frames like broadcast frames.
- In order to reduce the amount of entries in the CAM table, configure the network so that SW2 discards inactive entries from VLAN 8 and VLAN 88 after 10 seconds.

3 Points

2. IPv4

2.1 IGP Protocols

- Configure OSPF area 27 on the Ethernet segment between R2 and SW1.
- Advertise SW1's interface Loopback 0 into OSPF area 27.
- Configure your network so that the only OSPF route that SW1 sees is a default route generated by R2.
- Configure the EIGRP domain so that bandwidth, delay, and load are taken into account when computing metrics. The bandwidth should be three times more significant than either load or delay in the calculation.
- Configure RIPv2 on R1, R4, and R6. Enable RIP on VLAN 4, VLAN 6, VLAN 162, and the Frame Relay connection to BB1. Enable RIP on R6's interface Loopback 0.
- Configure R1 and R6 to authenticate all RIP updates received on VLAN162 with a secure hash value of the password CISCO. Use key ID 1 for this authentication.
- Configure R1 and R6 so that no unauthorized devices can *receive* RIP updates sent out on VLAN 162.

4 Points

2.2 IGP Features

- Configure the network in such a way that connectivity is maintained throughout the network if R5 loses its connection to the Frame Relay cloud. You are allowed to use static routes for backup purposes.
- Redistribute in the minimum number of places necessary to gain full reachability throughout the network.
- Routers in the OSPF domain should have the minimum amount of routing information needed to reach the RIP routes learned from BB3. Do not overlap unnecessary address space to accomplish this.
- Configure /16 summaries for the major network of your topology and for the loopback networks to advertise to BB3.
- **Note:** Networks that have not been added to an IGP (such as the networks interconnecting the switches) are not required to be reachable from the other devices in your topology.

4 Points

2.3 BGP Features

- Create a new Loopback interface on SW1 with the IP address 162.X.7.7/24 and advertise it into BGP. Create a new Loopback interface on SW2 with the IP address 162.X.18.8/24 and advertise it into BGP.
- Since SW1, SW2, SW3, and SW4 have connections only to AS 300, it has been decided that they will not apply for their own block of IP addresses, nor will they apply for a public BGP AS number. Instead, AS 300 has assigned them the locally significant AS numbers of 65001, 65002, and 65034.
- Configure your network so these AS numbers do not leak out onto the rest of the network when AS 300 is advertising prefixes that have been originated in AS 65001, AS 65002 or AS 65034. R1 through R6 should all be able to ping the above two loopback networks.
- Configure a new Loopback interface on R5 with the IP address 162.X.15.5/24 and advertise it into BGP. R4 should not pass this prefix on to any BGP speaking neighbors. The filtering configuration should be done on R5. Account for any other prefixes that AS 500 may advertise in the future.
- Configure R4 to add a penalty of 1000 to BGP prefixes each time a withdrawn message is received for them. R4 should stop advertising these unstable prefixes when their penalty value exceeds 3000. Once a stable prefix's penalty falls below 1000, it should be reinstalled in the BGP table as an active prefix. Ensure that no stable prefix's advertisement is withdrawn for more than 30 minutes.

7 Points

3. IPv6

Some IPv6 addressing has been pre-configured for this scenario.

3.1 IPv6 over Frame Relay

- Configure IPv6 on the Frame Relay link between R1 and R3 using the global unicast network 2001:CC1E:X:0::Y/64, where X and Y are decimal values for the rack and host number.
- Configure IPv6 on the Frame Relay links between R2, R3, and R4 using the site-local network FEC0:234::Y/64.
- Configure IPv6 BGP peering sessions between the following devices:

Device 1	Device 2
R4	R3
R3	R2
R3	R1

4 Points

3.2 IPv6 BGP Features

- Configure R1, R2, and R4 to advertise IPv6 networks of VLANs 162, 27, and 4 into BGP respectively.
- Configure R3 to advertise the IPv6 Frame Relay segments and VLAN 3 into BGP.
- Configure R3 so that R4 only sees one route to VLANs 3, 27, 162, and the Frame Relay link between R1 and R3.
- The advertisement should be as specific as possible while still encompassing all of these segments.
- R1 and R2 should not be affected by summarization.

4 Points

4. MPLS VPN

No scenarios in this section.

5. Multicast

Basic Multicast settings have been pre-configured for you. You need to discover the active multicast topology using IOS show commands.

5.1 RP Assignment

- Configure R3 and R5 to announce Loopback 0 interface as a candidate rendezvous-point (RP) through Auto-RP.
- For ease of management and future multicast configuration changes, configure R1 to map multicast groups 239.0.0.0 – 239.255.255.255 to R3 and multicast groups 226.0.0.0 – 238.255.255.255 to R5.
- Use the minimum number of access-lists and access-list entries on R1 to accomplish this.

3 Points

5.2 Multicast Features

- For security reasons, do not allow BB2 to become a PIM neighbor with R1 but allow R1 sending multicast traffic down on BB2 segment.
- Configure your network so that SW2 will not receive traffic for any administratively scoped multicast groups regardless of any IGMP join messages it receives for these groups.
- Configure the network so that multicast groups which use R3 as their RP never change to a shortest path source tree. Instead these multicast groups should always use a shared tree.

4 Points

6. Security

6.1 Traffic Filtering

- A new corporate policy has been put in to effect that requires R4 to secure its connection to BB3. R4 should treat its interface connecting to BB3 as an 'outside' interface and all other links as 'inside' interfaces.
- Any ICMP, UDP, or TCP traffic coming in from an inside interface and exiting the outside interface should be allowed to return.
- R4 should still allow all necessary routing protocol traffic in from the outside interface but deny any other traffic to the router.
- For management purposes, R4 will need to be able to ping and telnet to BB3.
- Do not use CBAC or ZFW to accomplish this task.

3 Points

6.2 DoS Prevention

- Recently, R1 and R6 underwent a ping DoS attack that originated from behind BB2. In response to this, your network administrator has requested you to configure R1 and R6 to not receive any ICMP echo requests sourced from the 205.90.31.0/24 network inbound on their interfaces attached to VLAN 162.
- Do not apply any configuration on either R1 or R6 to accomplish this.

3 Points

6.3 Packet Filtering

- Configure R6 to drop ICMP Echo packets with the size ranging from 100 to 200 bytes ingress on its Frame-Relay interface.
 - Only drop the packets leaving the router out of VLAN162 interface.
 - Do not use MQC or Flexible Packet Matching to accomplish this.

2 Points

6.4 Access Control

- Create three CLI roles on R4 named “SUPER”, “DEBUG”, “INTERFACE”
- The role “INTERFACE” should have access to all IP configuration commands for VLAN4 interface only.
- The role “DEBUG” should be able to use any `debug` and `undebug` commands.
- The last role name “SUPER” should encompass all other roles.
- Use the password of “CISCO” to authenticate all CLI roles.

3 Points

6.5 Control Plane Security

- Enhance R5’s protection by dropping packets going to all closed ports.
- Ensure this does not affect connections made on TCP ports 2020 and 2040.
- To reduce the effect of broadcast storms, limit the rate of input non-IP packets to 100 per second
- Limit the rate of ICMP packets to 10 per second.
- Ensure that all fragmented packets transiting the router are limited to 1 Mpps.

3 Points

7. Network Services

7.1 SNMP

- A new network management server has been installed to manage R6. Configure R6 using the following SNMP parameters:
 - Contact: CCIE Lab R6
 - Location: San Jose, CA US
 - Chassis ID: 556-123456
 - Read-Only community: CISCORO
 - Read-Write community: CISCORW
- The management station's IP address is 192.10.X.101.
- This is the only station that should be allowed to manage R6.
- Attempts by other devices to manage R6 via SNMP should be logged.
- The network management server will be expecting SNMP traps to use a community of CISCOTRAP and be sourced from R6's Loopback 0 interface.

3 Points

7.2 Syslog

- One of your network administrators has requested that R4 and R5 be configured to log all severity 5 and below messages to a syslog server with the IP address 192.10.X.101.
- This network administrator has configured the syslog server to expect these messages to use the SYS10 facility.
- R4 and R5 should include their hostname in the syslog messages.
- All syslog messages should be sourced from R4 and R5's Loopback 0 interfaces.

3 Points

7.3 DNS

- The network administrators have requested that they should be able to telnet to the routers in your network using their DNS names as opposed to their IP addresses while working on R6. The network administrator has setup a DNS server at IP address 192.10.X.100 for R6 to point to for DNS resolution.
- Ensure that if your administrators mistype a command when working on the console the router it does not try to resolve the mistyped command via DNS.
- This configuration should not affect any other lines on R6.

3 Points

7.4 Local Authorization

- The first level support engineers from the company's NOC have complained to management about not having access to view R6's running configuration.
- To appease them, configure R6 so that these users can see only the following information in the running configuration:
 - Hostname
 - Interfaces
 - Interface encapsulations
 - IP access-lists applied to interfaces
- The NOC users must enter privilege level 2 using the password CISCO prior to being able to view the configuration.

3 Points

7.5 WCCPv2

- Configure new WCCP service-group group in R5 using ID 99 and the multicast-group IP address "224.0.1.100". The cache-engines on VLAN2005 use the IP addresses of "162.X.5.100" and "162.X.5.200". Ensure no other engines could communicate with the router.
- Authenticate WCCPv2 communications using the password of "CISCO".
- Redirect the users coming in from VLAN55 interface; make sure no other subnets are redirected for this service.

3 Points

7.6 EEM

- Configure R6 so that every attempt to configure a Loopback interface settings is intercepted and the interface automatically placed in the shutdown state.
- This configuration should not affect any other interfaces but the Loopback interfaces.

3 Points

8. QoS

8.1 Frame Relay Traffic Shaping

- Configure Frame Relay Traffic Shaping on R1 per the following requirements:
 - R1 has a port speed of 512Kbps.
 - R1's DLCI 113 has a provisioned CIR of 256Kbps.
 - R1 should send data at 384Kbps and throttle down to CIR in the event of congestion notification from the Frame Relay cloud.
 - In the case that R1 has accumulated credit, it should be allowed to burst up to its port speed.
- Use an interval (Tc) of 100ms and don't use any legacy commands to accomplish this task.

2 Points

8.2 RTP Header Compression

- Configure the Frame Relay connection between R3 and R4 to support RTP header compression.
- This compression should support up to 15 connections.
- R3 should only compress RTP headers if it is receiving RTP headers that are compressed.
- R3 should not perform RTP header compression with any other routers.

2 Points

8.3 Bandwidth Limiting

- Users have been complaining about slow access to servers in VLAN 27. After further investigation, one of your network administrators has reported that the congestion appears to be caused by users accessing a Microsoft SQL server in that VLAN.
- To resolve this problem, configure your network so that Microsoft SQL traffic is limited to an average rate of 256Kbps on R2's connection to the Frame Relay cloud.
- Up to 2048 SQL packets in excess of 256Kbps should be queued up by R2 before packet loss occurs.
- Do not use an access-list to accomplish this.

3 Points

8.4 Catalyst QoS

- It has been discovered that SW1 and R6 are overwhelming R1 with IP and non IP-traffic.
- In order to alleviate this problem, configure SW1 to rate-limit traffic received on VLAN162 as following:
 - TCP traffic should be limited to 1Mbps and marked with CS0 and any packets exceeding this limit should be remarked to CS1.
 - UDP traffic should be limited to 512Kbps and any packets exceeding the limit should be dropped.
 - IPX packets received by SW1 on this VLAN should be marked with CoS value of 2 if they don't exceed 128Kbps and dropped if they do.
 - Classify IPX using the EtherType value of 0x8137.
- Do not apply the policy-map to any physical interface to accomplish this.

4 Points

IEWB-RS-VOL2 Lab 6

Difficulty Rating (10 highest): 7

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	14
IPv4	21
IPv6	8
MPLS VPN	7
Multicast	7
Security	6
Network Services	10
QoS	8

GOOD LUCK!

1. Layer 2 Technologies

1.1 Trunking

- Using 802.1Q encapsulation configure the following trunks:
 - SW1 Fa0/13 - SW2 Fa0/13
 - SW1 Fa0/16 - SW3 Fa0/13
 - SW1 Fa0/19 - SW4 Fa0/13
- SW1 should not trunk VLANs 7, 77, and 777 with SW3 and SW4.

3 Points

1.2 Spanning-Tree

- Ensure SW1 is forwarding on all trunk links for any active VLANs.
- If a new VLAN is added to the VTP domain NET12, SW1 should forward on all trunk links for the new VLAN.

3 Points

1.3 Layer 2 Tunneling

- Configure the network so that R4 and SW2 match the output below:

```
Rack1R4#show cdp neighbors fa0/1 | include SW2
Rack1SW2    Fas 0/1      121      S I      WS-C3560-2Fas 0/18

Rack1SW2#show cdp neighbors fa0/18 | include R4
Rack1R4     Fas 0/18     134      R S I    3640     Fas 0/1
```

- Use VLAN 100 if an additional VLAN is needed.
- Recabling of the network is not required for this task.

3 Points

1.4 MAC Filtering

- Port Fa0/10 of SW2 connects to an 802.11b wireless access point. Since there are only 4 hosts which should be accessing your network through this access point, the new corporate policy dictates that traffic from other hosts should not be allowed in this port. The MAC addresses of these four hosts are as follows:

Host	MAC Address
1	0050.7014.8ef0
2	00d0.586e.b710
3	00c0.144e.07bf
4	00d0.341c.7871

- Configure SW2 so that traffic is only allowed in this port if it is sourced from one of the above MAC addresses.
- In the case that other hosts try to access this port, a syslog message should be sent to the server 191.X.7.100.

3 Points

1.5 Spanning-Tree Convergence

- The Ethernet link connecting to the wireless access point has been periodically flapping and causing the wireless users to lose access to the network. After further investigation, you have determined that when the link comes up the users are subject to a delay as the spanning-tree process moves to the forwarding state.
- In order to minimize downtime for these users, configure SW2 so that this port goes immediately into the forwarding state when it is connected.
- As a precautionary measure, ensure that if a spanning-tree BPDU is received in this interface the normal forwarding delay is reinstated.
- Ensure that SW2 does not send any BPDUs on the mentioned interface until it has received a BPDU from the attached device.

2 Points

2. IPv4

Some IGP settings have been preconfigured for you.

2.1 OSPF

- Configure OSPF area 0 on the Frame Relay segment between R1, R2, and R5.
- Do not use the `ip ospf network` statement on any of these devices.

3 Points

2.2 OSPF Filtering

- Advertise SW1's Loopback0 interface into OSPF without assigning it to any area.
- Since SW1's only connection to the rest of the routing domain is through R2, it does not need specific routing information about the rest of your network.
- Configure the network so that R2 filters all routing advertisements to SW1 with the exception of a default route.
- Do not use a distribute-list or prefix-list to accomplish this.

3 Points

2.3 Conditional Default Routing

- R3 is the only connection between the OSPF domain and the other routing domains. In order to minimize the amount of memory necessary to maintain the routing table throughout the OSPF domain, configure your network so that all routers in the OSPF network send their traffic towards R3 if they do not have a longer match in their routing table.
- In order to prevent the unnecessary forwarding of traffic that will eventually be dropped, ensure that R3 **only** advertises this default route if it has an active connection to either BB2 or BB3.

3 Points

2.4 IGP Redistribution

- Redistribute VLAN 32 into RIP on R3.
- Redistribute between OSPF and RIP on R3.
- All routers in the OSPF domain should have a longer match for R6's interface Loopback 0.
- No other routes should be redistributed from RIP to OSPF.

3 Points

2.5 BGP Filtering

- Memory usage on your BGP speaking devices is getting dangerously high. After investigating the problem, you have determined that the BGP table is consuming too much memory. In order to help cut down on the memory requirements throughout the BGP domain, your design team has implemented a new filtering policy. This policy states that AS 100 will not accept any prefixes from AS 54 with a mask longer than a /20.
- Configure R6 to reflect this policy.
- The prefix-list used to accomplish this should only have one line.

2 Points

2.6 BGP Summarization

- Configure R3 to advertise a summary of your major network, 191.X.0.0/16, and your Loopback 0 addresses, 150.X.0.0/20, into BGP.
- Do not use the `aggregate-address` command to accomplish this.
- You are allowed to use two static routes on R3 to accomplish this.

3 Points

2.7 BGP Table Stability

- High CPU utilization has been reported on R6. After further investigation, you have discovered that the prefixes 112.0.0.0/8 and 113.0.0.0/8 from AS 54's customers have been constantly flapping and causing R6 to continuously recalculate the BGP topology.
- In order to minimize the impact of this flapping on the rest of the BGP domain, configure R6 so that these prefixes are not advertised if they are consistently unstable.
- No other prefixes should be affected by this configuration.

2 Points

3. IPv6

3.1 IPv6 Addressing

- The network administrator has requested you to configure a test deployment of IPv6 between VLAN 5 and BB2.
- Configure IPv6 on the Serial connection between R2 and R3, using the network 2001:CC1E:X:23::Y/64.
- Configure IPv6 on the Frame Relay connection between R1, R2 and R5, using the network 2001:CC1E:X:125::Y/64.
- Configure IPv6 on R5's interface connecting to VLAN 5, using the network 2001:CC1E:X:5::Y/64.

2 Points

3.2 RIPng

- Enable RIPng on all interfaces running IPv6.
- Create and advertise into RIPng three additional Loopback interfaces in R3 with the following IPv6 addresses:
 - 2001:220:20:3::1/64
 - 2001:222:22:2::1/64
 - 2001:205:90:31::1/64
- Configure R3 to originate a default route to R2 via RIPng.
- R2 should not see any of the above-mentioned IPv6 subnets
- Do not use a prefix-list and do not change metrics to accomplish this.

3 Points

3.3 EIGRPv6

- Configure the IPv6 subnet 2001:CC1E:X:45::/64 on the Ethernet link between R4 and R5.
- Create new Loopback interfaces in R4 with the following IPv6 addresses.
 - 2001:CC1E:X:444::4/64
 - 2001:CC1E:X:454::4/64
 - 2001:CC1E:X:484::4/64
- Configure IPv6 EIGRP AS100 between R4 and R5 and advertise a single optimal summary route for the above prefixes to R5.
- Redistribute between RIPng and EIGRPv6 to obtain full reachability.

3 Points

4. MPLS VPN

The ISP core network consisting of R4, R5 and R6 provide VPN services to the following customers:

Customer_A: BB1

Customer_B: SW3 and SW4.

VRFs, IP addressing and some basic IGP settings for these customers have been preconfigured for you. Some parts of MPLS VPN configuration might be missing and you should fill the gap yourself.

4.1 PE-CE Routing

- Configure EIGRP AS 10 for Customer_A on R6's connection to BB1.
- Administrators in your NOC have reported that R6 has been generating a "neighbor not on common subnet" log message for EIGRP. After further investigation, you have determined that a provisioning error on the part of your Frame Relay service provider is to blame.
- In order to avoid security issues with this type of problem in the future, configure R6 so that it does not accept any EIGRP packets on the Frame Relay interface except those sent from BB1.
- Additionally, authentication EIGRP packets using the key value of "CISCO".

3 Points

4.2 Backup Link

- Customer_B's sites are emulated by SW3 and SW4 use OSPF as the routing protocol. .
- Configure R4 and R5 so that the preferred path between SW3 and SW4 is across the link connecting R4 and R5.

4 Points

5. Multicast

Some of the multicast settings have been pre-configured for you. You need to discover the active multicast topology using the show commands.

5.1 PIM Filtering

- A media server located on VLAN 32 will be streaming a video feed to clients located on VLAN 5.
- The network administrator has requested that the Frame Relay connection between R1 and R5 be used as sparingly as possible for multicast traffic.
- To help avoid excess multicast flooding and pruning behavior over this Frame Relay connection, R1 should not allow R5 to become a PIM neighbor. However, R5 should still allow clients on VLAN 5 to receive multicast traffic for this group.
- Configure your network to support this arrangement.

3 Points

5.2 IGMP

- The network administrator has reported that clients in VLAN 363 will be using Windows® 95, which supports only IGMP version 1.
- Configure R3 to only support clients running IGMP version 1 on this interface.

2 Points

5.3 Multicast Testing

- The network administrator is trying to troubleshoot a problem relating to the multicast group 225.25.25.25 and has requested that SW1 forward traffic for this multicast group into VLAN 7. However, the testing application he is using will not be generating IGMP join messages.
- Configure SW1 to accommodate this request, but do not allow SW1 to process switch this traffic.

2 Points

6. Security

6.1 BPDU Filtering

- Recently, your managers brought in some outside consultants to perform a security audit of your network. After the audit these consultants have reported that there are unauthorized bridges in VLAN 363 sending DECnet spanning tree BPDUs. Until the source of these BPDUs can be located, the network administrator has requested that SW3 and SW4 filter off all DECnet spanning tree BPDUs in VLAN 363.
- Configure your network to accommodate this request.

3 Points

6.2 Traffic Filtering

- Recent network monitoring has shown a number of unauthorized sources attempting to telnet into SW1. In order to protect SW1 from unauthorized access, it has been decided to configure R2 to filter telnet traffic going to SW1.
- Configure the network in such a way that hosts must first authenticate to R2 before they are allowed to telnet to SW1.
- These users should authenticate with the username TELNET and the password CISCO.
- Users logging into R2 with the username CLI and password CISCO should be granted access to R2's CLI.
- Log user out after 5 minutes of inactivity.

3 Points

7. Network Services

7.1 SNMP

- Configure R3 to be managed via SNMP. R3 will be managed by two separate network management servers.
- The first network management server's IP address is 191.X.7.100 and second network management server's IP address is 191.X.77.100.
- The network management servers will be expecting SNMP traps to use community string CISCOTRAP.
- The network management servers will be expecting the RO community string to be CISCORO and the RW community string to be CISCORW.
- Only allow the first network management server to access the RW community string.
- Allow R3 to be reloaded via SNMP.

3 Points

7.2 RMON

- The network administrator is trying to do preventative maintenance by having R1 and R3 generate a log message whenever the utilization on the HDLC link between them exceeds twice the normal rate.
- The network administrator has determined that the average change of the input octet (ifEntry.10) value for R1 and R3's HDLC link is 40000 per minute.
- Configure R1 and R3 to generate a log message whenever this value reaches twice the average rate and again when it falls back below the average rate.
- R1 should monitor 'ifEntry.10.3' and R3 should monitor 'ifEntry.10.5'.
- The sampling interval should be every 60 seconds.
- The server to log these events to is 191.X.7.100.

3 Points

7.3 CDP

- One of your network administrators has written a custom network management application that relies on CDP to determine when a neighboring device is down, and would like to test this application on the Ethernet segment between R4 and SW2.
- For this application, waiting 60 seconds between sending CDP packets is too long. Configure R4 and SW2 to send CDP updates every 5 seconds.
- In addition to this, R4 and SW2 should discard a CDP entry if the neighbor has not sent a CDP update in over 15 seconds.
- The network administrator has also requested that all CDP packets sent by R4 include its Loopback 0 interface's IP address in the packet for identification.

2 Points

7.4 UDP Echo

- Configure SW2 to respond to UDP echoes from a network management station with the IP address 191.X.77.100.
- SW2 should not respond to packets sent to the UDP 'discard' and 'chargen' ports from this network management station.

2 Points

8. QoS

8.1 Real Time Protocol

- VoIP users connected to R4 have been complaining about poor voice quality. After further investigation, it has been determined that excessive HTTP traffic being sent over the Frame Relay connection between R3 and R4 is the likely cause.
- In order to resolve this problem, ensure that all RTP packets sent over the Frame Relay circuit between R3 and R4 are prioritized.
- Allocate 25% of the bandwidth for these RTP packets.
- This configuration should be done in such a way that it is easy to add additional QoS configuration at a later date.

3 Points

8.2 Congestion Avoidance

- Even after prioritizing RTP packets, your users are still having issues with low voice quality. In order to deal with this congestion, reserve 30% of the remaining interface bandwidth to all non HTTP traffic.
- Flows matching this class should be schedule using fair-queueing and use random early detection as queue management strategy.

3 Points

8.3 Link Optimization

- The link between R1 and R3 is provisioned at only 64Kbps.
- In order to decrease the interactive traffic latency, enable a technique that reduces TCP header overhead.
- Provide resources enough to support optimization for 32 bi-directional connections.

2 Points

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.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	12
IPv4	23
IPv6	5
MPLS VPN	0
Multicast	3
Security	8
Network Services	15
QoS	10

GOOD LUCK!

Troubleshooting

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

1. Layer 2 Technologies

Some Layer 2 settings have been pre-configured for you; VLANs have been created and assigned to the switchports according to the diagram supplied with the task.

1.1 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

- Configure the SW1 and SW4 in such a way that R3's Ethernet interface connected to SW1 and SW2 Fa0/20 appear directly connected via CDP.
- R1's Ethernet interface and SW2 Fa0/21 should also appear directly connect via CDP.
- If an additional VLANs is needed use VLANs 100 and 101.
- Configure SW1 Fa0/5 to limit unicast traffic inbound to 25% of the interfaces bandwidth.

4 Points

1.2 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. Ensure the IP Phones belong to the same VLAN.
- 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.

3 Points

1.3 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

1.4 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

2. IPv4

Some IPv4 IGP settings have been already pre-configured for you.

2.1 RIPv2

- Configure R6 so that it does not advertise its routes learned across the Frame Relay connection to either R2 or BB3.
- Configure both R2 and R6 to not accept any RIP advertisements from either BB1 or BB3.
- 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. Do not accept any RIP advertisements from BB2.
- The only route that should be advertised to BB2 through RIP is a summary of your internal address space 163.X.0.0/16 and the Loopback interfaces subnet 150.X.0.0/16. This summary should encompass your entire network and still be as specific as possible without unnecessarily overlapping address space.

3 Points

2.2 OSPF

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

```
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
```

- 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.
- Advertise SW3 and SW4's Loopback 0 interfaces via OSPF.
- These two networks should appear as below in SW1's routing table. Do not use redistribution to accomplish this task.

```
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.

5 Points

2.3 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.

4 Points

2.4 BGP Best-Path Selection

Basic BGP peering has been pre-configured for you. Discover the existing configuration using show commands and the diagram supplied with the scenario.

- 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.
- 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 the *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.

4 Points

2.5 BGP Features

- 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.
- 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.

3 Points

2.6 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 300 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.
- AS 300 also has multiple connections to AS 100. However due to the aggregation that you need to configure R1, 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.

4 Points

3. IPv6

3.1 IPv6 Features

- 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 R4 and 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 38 using the networks FEC0:CC1E:X:4::/64 and FEC0:CC1E:X:38/64 respectively.
- 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 OSPFv3 area 100 on all interfaces running IPv6.
- Traffic from VLAN 38 destined for the prefixes above should use the point-to-point Serial link between R4 and R5.
- If this link is down traffic from VLAN 38 to these prefixes should be rerouted over the Frame Relay circuit between R4 and R5.
- Configure R4 so that hosts running IPv6 in VLAN 38 do not have access to IPv6 enabled hosts in VLAN 4.
- Do not use a prefix-list to accomplish this.

5 Points

4. MPLS VPN

No scenarios in this section.

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 points assignments to accomplish this.

3 Points

6. Security

6.1 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

6.2 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 filter off all HTTP GET requests that contain 'root.exe' in them which come from BB2.

3 Points

6.3 Control Plane Policing

- Limit the aggregate rate of ARP traffic towards R1's route processor to 100 packets per second.
- The routing control traffic marked with an IP Precedence value of 6 should be limited to 50 packets per second.
- Drop all Telnet connections sourced from R5's VLAN5 IP address.
- Limit the rate of outgoing ICMP messages (e.g. ICMP responses) to 10 per second.

3 Points

7. Network Services

7.1 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 and SW2) should log using facility local5.
- These log messages should be time stamped with the current date and time, including the millisecond.

2 Points

7.2 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

7.3 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: INE.com

2 Points

7.4 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

7.5 Netflow

- Enable netflow statistics collection on R6's VLAN263 interface.
- Aggregate Netflow statistics in the router based on the source IP addresses up to /16 prefix length.
- Limit the aggregation cache size to 2048 entries and export Netflow statistics to the server at the IP address 163.X.6.100 using the UDP port number 9999
- Do not use Flexible Netflow to accomplish this and randomly sample every 3rd packet of a flow.
- Do not apply a flow-sampler map directly to the interface to accomplish this.

3 Points

7.6 Kron

- Configure a KRON occurrence on R4 that saves the running config to a remote TFTP server daily at 06:00.
- Use the server 163.X.4.100 and the filename "r4-config".

2 Points

7.7 Remote Command

- Configure the network so that R1 is allowed to view the running configuration of R6 using remote shell.
- Source remote-shell connections off the Loopback0 interface of R1
- Allow RCP file-copy from R6 to R1.

2 Points

8. 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 RTP traffic. 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.

8.1 Marking

- The first step in your end-to-end QoS policy is to ensure that all traffic is properly classified. 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.
- Ensure the traffic received by R4 and R5 on any other ports is marked with the DSCP value of CS1.

2 Points

8.2 Shaping

- The next portion of the QoS policy dictates that all traffic sent across the Frame Relay cloud should be shaped as 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 of R4 and R5 should be such as to minimize delay for the packet in shaper queue.
- 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

8.3 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

8.4 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

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.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	12
IPv4	22
IPv6	7
MPLS VPN	0
Multicast	10
Security	7
Network Services	9
QoS	12

GOOD LUCK!

1. Layer 2 Technologies

Some Layer 2 settings have been already preconfigured for you in this lab

1.1 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

1.2 Layer 2 Connectivity

- Configure SW2 and SW4 to allow communication for VLAN 26 between R2 and R6.
- Do not create VLAN 26 in any of the switches and do not create any additional VLANs to accomplish this.

2 Points

1.3 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

1.4 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

1.5 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.
- Use a Multilink group interface with the lowest number to accomplish this task.

3 Points

2. IPv4

Some IGP settings have been already preconfigured for you in this lab.

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 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

2.3 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

2.4 IGP Redistribution

- Redistribute between RIP and EIGRP on R5.
- Redistribute between OSPF and EIGRP on R1 and R3.
- Eliminate any IGP issues preventing full reachability.

3 Points

2.5 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

2.6 BGP Summarization

Confederation-based BGP peering mesh has been pre-configured for you. Discover the sub-confederations active BGP peering sessions using the relevant show commands and the diagram supplied with the scenario.

- 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

2.7 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` neighbor option to accomplish this.

2 Points

2.8 BGP Bestpath Selection

- Advertise VLANs 3, 4, and 7 into BGP.
- Configure the network in such a way that all traffic from AS 54 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

2.9 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

3. IPv6

IPv6 addressomg has been pre-configured on R1, R4 and R5. Use the respective show commands to discover additional information about the topology and IPv6 IGPs.

3.1 OSPFv3

- Configure OSPFv3 area 0 on the Ethernet segment of R1.
- Configure OSPFv3 area 0 on the Frame Relay segment between R1 and R5.

2 Points

3.2 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

3.3 IPv6 Redistribution

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

2 Points

4. MPLS VPN

No scenarios in this section.

5. IP Multicast

Multicast routing has been enabled on R1, R2, R3, R4 and R5. Protocol Independent Multicast has been enabled over the Frame-Relay clouds connecting R1, R4 and R5 as well as R2 and R3. Additional multicast path has been configure between R1 and R3.

5.1 Auto-RP

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

3 Points

5.2 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

5.3 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

5.4 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. Security

6.1 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 behind BB2 or BB3.
 - Drop all inbound echo requests coming from behind BB2 or BB3.

3 Points

6.2 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

6.3 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

7. Network Services

7.1 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

7.2 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

7.3 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

7.4 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.
- Use the protocol that does not utilize the multicast address 224.0.0.2 and runs over UDP.

3 Points

8. QoS

8.1 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.
- Do not use the `map-class frame-relay` command to accomplish this.

3 Points

8.2 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.

3 Points

8.3 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

8.4 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

IEWB-RS-VOL2 Lab 9

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	7
IPv4	20
IPv6	4
MPLS VPN	6
IP Multicast	6
Security	10
Network Services	14
QoS	12

GOOD LUCK!

1. Layer 2 Technologies

1.1 Layer 2 Features

- Recently, the administrators in your NOC have reported strange traffic patterns throughout your switch block. After further investigation, you have discovered that one of your customer's switches connected to port Fa0/24 on SW2 has been advertising superior BPDUs into your network. After talking to your customer's engineers, the problem has been resolved, but your management is concerned about this happening again in the future.
- Configure SW2 so that port Fa0/24 is disabled if superior BPDUs are received.
- Furthermore, configure SW2 to match the highlighted command output below:

```
Rack1SW2#show spanning-tree vlan 68
```

```
VLAN0068
  Spanning tree enabled protocol ieee
  Root ID    Priority    24644
             Address    0016.9d31.8380
             This bridge is the root
             Hello Time 1 sec Max Age 7 sec Forward Delay 5
sec
```

- Use the fewest commands needed to accomplish the above show command output in SW2.
- Next, configure a logical layer 3 Etherchannel link between SW3 and SW4 using the Fa0/14 and Fa0/15 connect to SW1 on each switch; there should be no negotiation involved.
- Make sure to match the highlighted command output below:

```
Rack1SW3#show cdp neighbors | i 0/14|0/15
```

```
Rack1SW4    Fas 0/15    151        R S I    WS-C3560- Fas 0/15
Rack1SW4    Fas 0/14    151        R S I    WS-C3560- Fas 0/14
```

- Use the IP addressing and PortChannel number from the diagram.
- If additional VLANs are needed, you may use VLANs 100 and 101.

4 Points

1.2 WAN Links

- Configure the Frame Relay connection between R3 and R5.
 - Do not use subinterfaces on either R3 or R5.
 - Do not use the `frame-relay map` command on either R3 or R5.
 - Do not allow Inverse-ARP requests to be sent out any DLCIs other than 315 and 513.

- A point-to-point Serial link has been provisioned between R4 and R5 in order to maintain connectivity in the case that R4 loses its connection to the Frame Relay cloud.
 - Configure the network so that if the line protocol of R4's subinterface goes down this interface becomes active.
 - Once R4 regains its connection to the Frame Relay cloud, it should wait for 5 minutes before shutting the Serial link down.

3 Points

2. IPv4

2.1 OSPF

- Configure R1 and SW2 to stop generating log messages when an OSPF type 6 LSA is received.
- In order to ensure that false routing information is not injected into the OSPF domain from VLAN 18, configure R1 and SW2 to authenticate all adjacencies established on this segment.
- Do not use the `ip ospf authentication message-digest` command to accomplish this.
- R1 and SW2 should use the MD5 key number 7 with the string of CISCO, and ensure that all passwords are stored in both devices' configuration in an encrypted form.
- Ensure that SW2 does not authenticate the OSPF neighbor relationship with R6.

3 Points

2.2 EIGRP

- In order to help the network converge faster in the event of the Frame Relay link failing, configure R3 and R5 to declare their neighbor relationship dead if they have not received an EIGRP hello in 12 seconds.
- The network administrator has requested that R5 and SW1 authenticate each other using the password of CISCO2010.
- For added security, R5 and SW1 should rotate their keys used for this authentication. This key rotation should occur at 11:45 PM Dec 31st, 2010.
- The new key to use for authentication is CISCO2011.
- To help ensure that R5 and SW1's key rotation does not result in a network outage, allow for the both keys to be accepted 30 minutes prior to and after the scheduled key rotation time.

3 Points

2.3 RIP

- The network administrator has noticed that BB2 is sending RIP updates into VLAN 232. After several failed attempts to contact the team responsible for managing BB2, your network administrator has requested that R2 and R3 not receive any RIP updates sourced from BB2.
- This configuration should not be performed on R2 or R3.
- Do not use any deny statements in your access-lists.

2 Points

2.4 Network Migration

- Without removing or altering the EIGRP configuration, migrate SW3 and SW4 to RIPv2.
- SW3 and SW4 should use RIPv2 for reachability to the rest of the network, but the rest of the network should use EIGRP for reachability to the Ethernet segment between them and their Loopback0 subnets.
- SW3 and SW4 should use EIGRP routes for reachability to subnets within the 150.X.0.0/16 network.
- R3 and SW1 should run RIPv2 with SW3 and SW4 respectively, along with EIGRP.
- Ensure SW1 does not send routing updates out of non-transit interfaces.

3 Points

2.5 IGP Redistribution

- Redistribute between RIP and OSPF on R1.
- Redistribute between RIP and EIGRP on R3.
- When R4's connection to the Frame Relay cloud is down, the only IGP route it should see is a default route pointing to R5.
- Ensure that all devices have connectivity to R4's Serial and Loopback 0 networks when its Frame Relay connection is down.
- You are allowed one static route to accomplish this.
- Redistribute between RIP and EIGRP on SW1.
- Configure SW1 so that BB3 cannot advertise routes learned from SW1 onto other RIP speaking devices. Directly connected subnets of SW1 are excluded from this requirement.
 - Do not use an `offset-list` to accomplish this.

3 Points

2.6 BGP Features

- In order to reduce the amount of memory required by the BGP process, configure R6 to only accept prefixes from BB1 that have been originated by themselves and their directly connected customers.
- Configure R4 to advertise the 10.X.4.0/24 subnet into BGP. Configure R1's unused Ethernet interface to accomplish this.
 - Ensure that the 10.X.4.0/24 prefix shows up on R1 as 10.0.0.0/8.
 - Do not use the **aggregate address** or **network** command to accomplish this task.

3 Points

2.7 BGP Summarization

- Configure R6 to advertise an aggregate of your internal address space as well as the 54.X.3.0/24 subnet into BGP.
- Since the Frame Relay link is AS 54's only connection to your network, it does not need specific subnet information about your address space.
- Configure your network so that BB1 has the minimum amount of information necessary to obtain reachability to the 148.x networks of your topology.
- Do not use either the **default-originate** or **summary-only** keywords to accomplish this.
- Create an additional Loopback interface on SW1 using the 148.X.177.0/24 subnet and advertise it into BGP.
- This prefix should not be advertised outside of AS 65057.
- Ensure that R5 still has reachability to this network.
- The filtering configuration should be done on SW1.

3 Points

3. IPv6

3.1 IPv6 Features

- Hosts on VLANs 3, 4, 5, and 6 need to communicate with each other via IPv6, however, you don't want to enable IPv6 on every device in the transit path between these devices. In addition to this, you do not want to have to maintain manual point-to-point tunnel configurations as more IPv6 enabled segments come on to your network.
- Configure R3, R4, R5, and R6 in such a way to allow fully meshed connectivity between their IPv6 enabled VLANs.
- IPv6 addresses used should be in the format 2002:96XX:Y0Y:3456::Y/64, where Y is the router number in decimal notation and XX is the rack number in *hexadecimal* e.g. 0A for rack10 or 20 for rack 32.
- This configuration should dynamically account for new IPv6 enabled segments being added in the future.
- You are allowed one non-default static IPv6 route on each of these devices to accomplish this.
- Recent network monitoring has indicated numerous failed attempts to telnet to R6 via IPv6.
- In order to prevent unauthorized access to R6, configure the network so that only the management station PC is allowed to start telnet sessions to the command line of R6.
- Do not use the `ipv6 traffic-filter` command to accomplish this.
- The management station PC is located on VLAN 6 and has a host address of 0209:6BFF:FE06:47EF.

4 Points

4. MPLS VPN

4.1 VRFs

- Create two new loopback interfaces on R4:
 - Loopback 40 – IP address 40.40.40.4 255.255.255.0
 - Loopback 41 – IP address 41.41.41.4 255.255.255.0
- Create the following new loopback interfaces on R5:
 - Loopback 50 – IP address 50.50.50.5 255.255.255.0
 - Loopback 51 – IP address 51.51.51.5 255.255.255.0
- Create a VRF on R4 named R4TEST. Create a VRF on R5 named R5TEST.
- Add the two new loopbacks on each device to the VRF.

3 Points

4.2 VRF Routing

- Configure a new logical interface between R4 and R5, using the address space 148.x.45.4/31. Add the logical interface to the VRFs on R4 and R5.
- Configure R4 and R5 for EIGRP as the routing protocol for the VRFs, using 45 as the autonomous system number. Do not use the command `router eigrp 45`.
- Ensure that R4 and R5 can each ping the VRF loopbacks of the other device.

3 Points

5. IP Multicast

5.1 Multicast Testing

- A Windows media server located on VLAN 3 is streaming a multicast video feed into your network. You have received complaints from users in VLAN 6 that they are unable to receive these feeds.
- Configure the network to resolve this problem.
- For further testing purposes, ensure that R6 responds to ICMP echo requests sent to the multicast group 224.6.6.6 sourced from VLAN 3.
- Do not change PIM mode settings on any interface.

2 Points

5.2 Multicast RPF

- Your NOC engineers have reported excessively high CPU utilization on R2. While investigating the problem, you have noticed that various unstable unicast routes are causing excessive amounts of triggered RPF checks.
- In order to help alleviate this problem, configure R2 so that it waits at least 300ms between consecutive RPF checks.

2 Points

5.3 IGMP

- Configure the designated IGMP querier on VLAN68 segment so that failed multicast traffic receivers are detected and removed within 60 seconds.
- Every active receiver should respond to general IGMP queries within 16 seconds.
- Designated querier failures should be detected 3 times faster than by default.

2 Points

6. Security

6.1 DoS Filtering

- Recently, the administrators in your NOC have notified you that an excessive number of ICMP packets are being received on the Frame Relay link to the Internet. After further investigation, you have determined that you are undergoing a DoS attack which is originating from spoofed private addresses.
- In order to reduce the impact of this attack on your internal network, configure R6 so that it does not accept traffic from the Internet if it is sourced from these hosts as defined in RFC 1918.

2 Points

6.2 Traffic Filtering

- Recently, application monitoring has shown that users on VLAN 5 have been excessively surfing the Internet during work hours. In response to this, your manager has requested that you configure R5 to block these users' activities so that they can only go to your internal web server at 148.X.3.100.
- After work hours, these users should be allowed full access.
- Work hours are from 9 AM to 5 PM Monday through Friday.

2 Points

6.3 Access Control

- Only allow Telnet connections to R2 from the Loopback0 subnets.
- Create a new local user in R2 and authenticate it locally using the name "TELNET" and the password of "CISCO".
- Once logged in, this user should be allowed to connect to R3 only.

3 Points

6.4 Control Plane Security

- Configure R6's Frame-Relay respond interface not to respond with an ICMP message when the router drops an ingress packet.
- The rate of ICMP unreachable messages sent out of all other interfaces should not exceed 100 per second.
- In order to ensure correct operations of PMTUD, increase the rate of ICMP messages used by MTU discover process to 1000 per second.

3 Points

7. Network Services

7.1 Crash Logging

- One of your network administrators has reported that R6 has been experiencing random crashes. After consulting with TAC, they have recommended that a core dump be captured from R6 if it crashes again.
- Configure R6 to send a core dump via FTP to the server 148.X.3.100.
- The file name to send is R6DUMP.txt.
- Use the username R6CORE and the password CISCO when sending this file to the FTP server.

3 Points

7.2 NTP

- Recently, there was a brief network outage due to a misconfiguration in the EIGRP authentication between R5 and SW1. After further investigation, you have verified that the configuration was correct, but it appears that the system clocks were not consistent between R5 and SW1. In order to prevent this problem in the future, you have decided to implement Network Time Protocol on R5 and SW1.
- Configure R5 and SW1 to get network time from BB3.
- In the case that BB3 is unreachable, R5 and SW1 should be able to maintain consistent time amongst them.

2 Points

7.3 NTP Authentication

- To ensure the legitimacy of their time sources, configure R5 and SW1 to authenticate the NTP information coming from BB3 using an MD5 hash of the password CISCO and key index of 1.

2 Points

7.4 TCP Session Establishment

- While telnetting to one of your network devices from R1, you accidentally mistyped the IP address and were forced to wait 30 second for the router to return to the CLI prompt.
- In order to avoid this long delay, configure R1 to cancel a TCP request if the session has not reached the established state within 5 seconds.

2 Points

7.5 Traffic Monitoring

- For capacity planning purposes, your manager would like to know which hosts are sending the most traffic out the Frame Relay link to BB1.
- Configure R6 to collect these statistics for your manager and store them locally.
- To ensure that this configuration does not negatively impact your network, do not allow R6 to store more than 1000 entries.

2 Points

7.6 NAT Load Balancing

- Recent utilization monitoring on your internal web server has shown that it is becoming overloaded with HTTP requests. In order to alleviate congestion and speed up response time, three new servers have been installed on VLAN 3.
- Configure NAT on R3 so that traffic is transparently load balanced between these new servers without having to inform the users of the server change.
- The old web server's address was 148.X.3.100.
- The new server addresses are 148.X.3.110, 148.X.3.111, and 148.X.3.112.
- These servers support web requests at ports 80, 443, and 8080.

3 Points

8. QoS

8.1 Frame Relay Traffic Shaping

- Your company has recently purchased a 5Mbps Internet connection from the Frame Relay provider to BB1. However, the lowest speed interface that the provider supports to accommodate this connection is DS3.
- To prevent the dropping of your traffic, configure R6 network so that traffic sent out to BB1 does not exceed 5Mbps on average.
- The provider has agreed to allow you to burst up to 7.5Mbps for a maximum period of 32ms.
- Do not use the `frame-relay traffic-shaping` command to accomplish this.

3 Points

8.2 Policing

- After implementing traffic shaping, the help desk has been getting a lot of complaints about slow network performance. After further investigation, it appears that someone inside your network is sharing files through a peer-to-peer file sharing application. Instead of blocking this traffic, your design team has suggested that you police this type of traffic to the lowest values possible. Therefore, users attempting to download files from your network will become frustrated and give up.
- Ensure to include KaZaA, Morpheus, BearShare, and LimeWire traffic in this policy.

3 Points

8.3 Congestion Management

- Even after implementing the above policy, your administrators have still been getting numerous complaints from users about slow network response time. The majority of these users are complaining that it is taking a very long time to send e-mail and access the web.
- In order to increase performance for these users, configure R6 so that HTTP traffic is guaranteed a minimum of 2Mbps of the output queue on the Frame Relay link, while SMTP traffic is guaranteed a minimum of 1Mbps of the output queue.

2 Points

8.4 Congestion Management

- After the last addition to your QoS policy, your administrators have reported that the number of complaints from network users has dropped dramatically. However, now you have noticed that the ping time to #ccie on irc.internetworkexpert.com is horribly slow.
- In order to decrease your latency to the channel, configure R6 so that up to 32Kbps of your IRC traffic (TCP 6667) is dequeued first out the Frame Relay link to the Internet.
- Your PC's IP address is 148.X.6.10.

2 Points

8.5 Selective Packet Discard

- Enable aggressive Selective Packet Discard for ingress queues on R3.
- The headroom for IP precedence 6 packets should be set to 50 packets.
- The headroom for incoming HSRP packets should be set to 100 packets.
- Start dropping low-priority packets randomly when the input queue is 50% full.

2 Points

IEWB-RS-VOL2 Lab 10

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 internal to 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.
- Do not configure any other device but R3 to accomplish this.

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 Serial 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 OSPFv3.
- 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.
- Do not apply an access-group to the interface to accomplish this.

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.

```
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 it 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 VoIP latency on the circuits, ensure that all the shaping intervals are the smallest possible, and that a single packet cannot take more than 10ms to be serialized.

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.
- Classify VoIP bearer traffic base on UDP port range 16384-32767.
- Do not use the `ip rtp priority` command to accomplish this.

3 Points

IEWB-RS-VOL2 Lab 11

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	16
IPv4	15
IPv6	7
MPLS Technologies	6
Multicast	11
Security	6
Network Services	6
QoS	12

GOOD LUCK!

1. Layer 2 Technologies

Basic VTP settings and VLANs are preconfigured for this lab.

1.1 STP Features

- Configure SW1 as the spanning-tree root for all even numbered VLANs.
- Configure SW2 as the spanning-tree root for all odd numbered VLANs.
- Do not make any changes to SW2 configuration to implement the above requirements.
- Configure SW1 so that all even numbered VLANs prefer the Fa0/14 trunk link over the Fa0/13 and Fa0/15 trunks.
- In the event of Fa0/14's failure, all even numbered VLANs should switch over to the Fa0/15 trunk.
- Even numbered VLANs should only use the Fa0/13 trunk in the event that both Fa0/14 and Fa0/15 fail.
- Do not apply any configuration commands on SW1's interface Fa0/13 or any interface of SW2 in order to accomplish the above requirements..
- Configure SW1 so that all odd numbered VLANs prefer the Fa0/15 trunk link over the Fa0/13 and Fa0/14 trunks.
- In the event of Fa0/15's failure, all odd numbered VLANs should switch over to the Fa0/13 trunk.
- Odd numbered VLANs should only use the Fa0/14 trunk in the event that both Fa0/13 and Fa0/15 fail.
- Do not apply any configuration commands on SW1's interface Fa0/14 or any interface of SW2 in order to accomplish the above requirements.

4 Points

1.2 MAC Address Notifications

- A recent security breach which involved the compromising of the company's future business plans was tracked down to a notebook computer that was located in VLAN 28 with a MAC address of 0001.02ac.9ab2. After checking the MAC address tables of SW1 and SW2, you have determined that the notebook computer is not currently plugged into the network.
- In order to help track down this device in the future, configure SW2 to notify the network management station at 187.X.3.100 whenever a new MAC address is learned in VLAN 28.
- The network management server will be expecting community-string to be CISCOTRAP.

3 Points

1.3 Layer 2 Features

- After numerous attempts to get the company's graphics department to migrate their legacy servers to IP, you have decided to configure the network to only allow IPv4 traffic and necessary layer 2 traffic to transit VLAN 56.
- Use a named filter called IPONLY to accomplish this.
- The classic STP protocol is known for its slow response time to network failures.
- In order to improve network convergence, configure the switches so that any device is able to detect and respond to an indirect link failure in 20 seconds.

3 Points

1.4 IP Phones

- Interfaces Fa0/7 and Fa0/8 on SW1 connect to Cisco 7960 IP phones.
- VoIP originating from these phones is being marked with a CoS of 5.
- This VoIP traffic should belong to VLAN 7.
- Traffic coming from the PCs connected to the access ports of these IP phones should belong to VLAN 17.
- Ensure that all traffic originating from the IP phones maintains its CoS values while transiting your switched network, while traffic coming from their attached PCs is set to 0.
- For ease in future changes of these interfaces, configure SW1 so that these ports can be configured at the same time by using a macro named VPORTS.

3 Points

1.5 PPP Authentication

- Configure PPP on the Serial link between R4 and R5.
- R4 should challenge R5 to authenticate using CHAP.
- R5 should respond with the username RackXR5 and the password C1SC0?2000.
- Use an imaginary RADIUS server at the IP address 192.10.X.100 authenticated using the key value of CISCO.
- Communicate with the server using the Loopback0 interface to source the UDP packets and declare the server dead for a minute of no response.
- Fall back to the local authentication if the RADIUS server fails.
- Do not use the `username` command on R5 to accomplish this.

3 Points

2. IPv4

Some IGP settings and IP addressing have been pre-configured for this lab. Use the diagram as your reference for more information.

2.1 OSPF

- Configure OSPF area 134 on the Frame Relay cloud between R1, R3, and R4.
- Configure this OSPF network in such a way that R1 sees OSPF routes advertised by R4 with a next hop value of R3, and vice versa.
- SW2 should advertise its Loopback in area 38.
- SW1 should advertise its Loopback in area 0 and other prefixes per the diagram supplied
- Ensure that R5, SW1, and SW2 see this Frame Relay subnet as 187.X.134.0/24.
- Using the password of CISCO, authenticate the OSPF virtual link between R3 and R4 using the strongest authentication method supported by OSPF.
- Do not authenticate any other virtual links using this method.
- Using the password of CISCO, authenticate the OSPF virtual link between R1 and R3 using simple password authentication.
- Make sure the adjacency between R1 and SW1 is never authenticated.

4 Points

2.2 IGP Features

- Recent network monitoring has shown an excessive exchange of EIGRP query messages between R2, R3, and R5 when a route in the EIGRP domain is lost.
- Configure the network in such a way that EIGRP query messages are not sent to R2 in the event of a network failure anywhere in the EIGRP domain.
- Ensure that connectivity remains throughout the EIGRP domain if one of the circuits between R2, R3, and R5 goes down.
- Configure RIP on SW2 as per the diagram and redistribute OSPF into RIP on SW2. SW2 should use MD5 with a string of "CISCO" and key number 1 to authenticate RIP updates with BB2.
- Redistribute between EIGRP and OSPF on R3 and R5.
- EIGRP and RIP routes with an even numbered first octet should be redistributed into OSPF as E1; Odd routes should appear as E2 with a metric of 100.

3 Points

2.3 BGP Features

BGP peering sessions have been preconfigured for this lab. Refer to the diagram and initial configurations for more information.

- Create a new Loopback interface on SW1 with the IP address 187.X.77.7/24 and advertise it into BGP.
- From the perspective of BGP speaking devices beyond AS 200, this prefix should appear to have originated in AS 200.
- Advertise the Frame Relay subnet between R2, R3, and R5 (187.X.235.0/24) into BGP.
- Configure R2 and R6 to advertise a single route representing your entire primary network (187.X.0.0/16) to BB1 and BB3.
- To ensure that AS 54 uses R2 as the entry point for the 187.X.235.0/24 prefix, configure R2 to continue sending the specific route 187.X.235.0/24 along the aggregate 187.X.0.0/16.

4 Points

2.4 BGP Advertisements

- Configure SW2 to advertise its interface VLAN 28 along with all routes learned via RIP from BB2 into BGP.
- Do not use redistribution or aggregation to accomplish this.
- AS 100 provides the following services to customers:
 - Prefixes tagged with BGP community 100:542 should be advertised to BB3 with AS 100 prepended 3 times
 - Prefixes tagged with BGP community 100:546 should be advertised to BB1 with AS 100 prepended 3 times
- To test this configuration, create two Loopback interfaces in R3 with IP addresses 150.X.33.33/24 and 150.X.133.133/24.
- Using the above developed signaling, ensure that AS 54 prefers BB1 link to reach 150.X.33.33/24 and prefers BB3 link to reach 150.X.133.133/24

4 Points

3. IPv6

3.1 IPv6 Addressing & Routing

- Configure IPv6 on VLAN 4 of R4 using the network 2001:187:X:4::/64.
- Configure IPv6 on VLAN 17 of R1 and SW1 using the network 2001:187:X:17::/64.
- Configure IPv6 on VLAN 56 of R6 using the network 2001:187:X:56::/64.
- Configure fully meshed IPv6 over IPv4 tunnels between R1, R4, and R6.
- Use the default encapsulation for these tunnels, and addressing in the format 2001:187:X:AB::/64 where “A” is the lower of the routers’ numbers and “B” is the higher.
- Enable OSPFv3 on VLANs 4, 17, 56 of R1, R4, R6 and SW1 as well as on the IPv6 over IPv4 tunnels, using instance 99.
- R1 should advertise to SW1 only a single summary prefix encompassing IPv6 subnets of VLANs 4 and 56.

4 Points

3.2 Traffic Engineering

- Configure the network in such a way that IPv6 traffic from VLAN 17 going to VLAN 56 is first sent to R4, and then on to R6.
- Traffic from VLAN 56 back to R1 should be sent directly from R6 to R1.
- Do not use summarization techniques to accomplish this.

3 Points

4. MPLS

Note: In the initial configuration file, SW4 has been preconfigured for VRFs for the connections to R4 and R6, and will be acting as the CEs. Verify that you have correctly loaded the initial configuration for SW4.

4.1 VRF

- Configure R4 for a VRF named R4. Add the Fa0/1 interface to this VRF, and configure for RIP.
- R4 should receive routes from SW4.
- Configure R6 for a VRF named R6. Add the Fa0/1 interface to this VRF, and configure for EIGRP AS 99.
- R6 should receive routes from SW4.

3 Points

4.2 VPNv4 BGP

- Configure a VPNv4 BGP peering between R4 and R6. Configure Redistribution to/from BGP on R4 and R6 for the respective BGP IPv4 address families.
- Verify that SW4 receives routes from the other side, and can ping them when sourcing from the loopback interfaces. You may add two addresses for this section.

3 Points

5. IP Multicast

IP multicast routing is configured on R1, R3, R4, R5, and SW1. PIM sparse mode is enabled on VLANs 3, 4, 5, 7, and 17. Additionally, PIM sparse mode is enabled the Frame Relay segments between R1, R3, & R4, and R2, R3, & R5.

5.1 RP Advertisement

- Configure R4 to announce itself as the RP for the multicast groups 224.0.0.0 – 231.255.255.255.
- Configure R5 to announce itself as the RP for the multicast groups 232.0.0.0 – 239.255.255.255.
- R3 should be responsible for group to RP mappings.
- Do not use Auto-RP to accomplish this.

3 Points

5.2 Multicast Testing

- One of your network administrators has informed you that his PC in VLAN 7 is unable to receive the multicast feed 228.34.28.100 that is being originated from a server in VLAN 4.
- Configure the network to resolve this problem, and so that SW1 responds to ICMP echo requests sent to 228.34.28.100 coming from VLAN 4.
- Do not use the `ip pim nbma-mode` command to accomplish this.

2 Points

5.3 Multicast Security

- Configure the RPs so that only sources connected to R3 and R4 are allowed to send multicast traffic streams.
- All multicast-enabled routers should permit building of multicast shared trees towards the two configured RPs only.

3 Points

5.4 Layer 2 Multicast

- Configure SW4 so that R4 and R6 interfaces are in VLANs 44 and 66 respectively.
- Without enabling multicast routing on the path between R4 and R6 ensure multicast feeds from R4 are delivered to R6.
- To test your configuration, join R6's VLAN66 to the multicast group 239.6.6.6 and make sure R4 can ping it.
- R4 should not flood multicast traffic unless there are no receivers on VLAN66.

3 Points

6. Security

6.1 Protocol Filtering

- Recently a CERT security advisory was released that reported various vulnerabilities in the version of IOS used in your network. In response to this, Cisco has recommended that IP protocols 53, 55, 77, and 103 be denied from both entering and leaving the network.
- Configure a filtering policy on R2, R6, and SW2 to reflect these new recommendations.
- In order to minimize the impact of this filtering policy on these devices, ensure that TCP and UDP traffic is permitted prior to denying any other IP protocols.
- Your security team has expressed interested in the amount of packets that are denied by this filtering policy and have requested that denied packets be logged to a syslog server at 187.X.38.100.
- Configure these devices to reflect this requirement.

3 Points

6.2 IOS Firewall

- Your security manager requests to implement a security policy on R6's connection with BB1 per the requirements below:
 - Only allow packets ingress on the connection to BB1 if they are part of an already established session.
 - Allow outgoing HTTP, FTP and DNS sessions.
 - Collect information on the number of bytes transferred with every HTTP and FTP session.
 - Allow for VoIP calls originate from R6 to H.323 gateways behind BB1.
- Make sure that your configuration does not affect any routing protocol configured.

3 Points

7. Network Services

7.1 Change Logging

- Recently, users were unable to access resources from BB1. This was due to the fact that one of your administrators misconfigured an access-list on R6. Unfortunately, you are not sure which admin it was since logging wasn't enabled.
- To avoid this problem in the future, implement a change control policy on R6 which logs all commands entered to syslog.
- The syslog server's IP address is 187.X.5.155.
- In the case that the syslog server is unavailable, R6 should store up to 500 of these log entries locally.

3 Points

7.2 Web Acceleration

- Due to the large amount of time that some of your coworkers spend browsing the Internet, you have recommended to management that a web cache engine be installed to enhance their Internet browsing experience. As usual, management has blindly taken your recommendation and approves the purchase of a web cache engine. Your coworkers that will need to have their HTTP requests redirected to the web cache engine are located in VLAN 3.
- Your personal Linux workstation is also located in VLAN 3. Since you do not have time to browse the Internet like some of your coworkers, you have decided to exclude your HTTP requests from being cached.
- Your workstation's IP address is 187.X.3.50.
- Configure R3 to reflect this policy.

3 Points

8. QoS

8.1 Traffic Shaping

- After recent connectivity issues between R1, R3, and R4, you have noticed that a large percentage of frames arriving from R3 have the DE bit set. After discussing this issue with the Frame Relay service provider's helpdesk, they have recommended that Frame Relay Traffic Shaping be enabled on R3.
- Configuring FRTS on R3 according to the following parameters:
 - R3's connection to the Frame Relay cloud has a port speed of 512Kbps.
 - A CIR of 192Kbps was subscribed with the Frame Relay service provider for DLCI 301 and 304.
 - Allow either DLCI to burst above CIR if credit is available.
- To help ensure that one DLCI does not ever consume all the bandwidth, only allow bursts up to 320Kbps for a maximum period of 100ms.

3 Points

8.2 Traffic Marking

- R1's connection to the Frame-Relay cloud is oversubscribed, and the router is allowed to send up to the physical access rate.
- To reduce the possibility of VoIP and critical network traffic being dropped in Frame-Relay cloud, configure R1 to mark all outgoing packets except VoIP and OSPF packets with DE-bit.
- Classify VoIP packets based on their size between 80 and 100 bytes.
- Do not use any legacy commands to accomplish this task.

3 Points

8.3 Rate-Limiting

- BB2's connection to SW2 is 10Mbps but the contracted output rate is 3Mbps.
- Configure SW2 to enforce this requirement and ensure that packets exceeding the contract are not dropped.
- Packets marked with DSCP value of AF31 should be limited to 1Mbps of output rate.
- Use the default QoS mapping tables to accomplish this task.

3 Points

8.4 Traffic Classification

- R5 connects two customers behind VLAN5 and VLAN56 to the Frame-Relay service provider.
- The Frame-Relay contracted connection rate is 512Kbps and equals R5's physical connection rate.
- Each customer is guaranteed 256Kbps of the WAN link rate, but may send up to the maximum port speed.
- Configure R5 to meter incoming traffic rate on both Ethernet connections and mark traffic conforming to the contract with IP Precedence value of 1.
- At the same time, traffic above the contracted rate should be marked with IP Precedence of 0. All traffic received above the WAN port egress physical rate should be dropped.

3 Points

IEWB-RS-VOL2 Lab 12

Difficulty Rating (10 highest): 7

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	9
IPv4	22
IPv6	10
MPLS VPN	8
Multicast	6
Security	5
Network Services	10
QoS	9

GOOD LUCK!

1. Layer 2 Technologies

1.1 MAC Filtering

- SW1's interfaces Fa0/7 and Fa0/8 are connected to the company's public meeting room. Your corporate policy dictates that these ports should not be connected to a hub or switch to split the connection, however, your users have not been cooperating.
- In order to limit the number of PCs that can connect to the network through these ports, configure SW1 to shutdown an interface connected to the meeting room for 60 seconds if it learns more than two MAC addresses on it.
- Additionally, you have suspicions that a sales engineer has plans to circumvent the two MAC address limitation in the meeting room by connecting a router to one of the RJ-45 jacks in the room.
- Configure SW1 to stop this router which has the MAC address of 0030.1369.87a0 from communicating if it is connected to either interface Fa0/7 or Fa0/8.

3 Points

1.2 QoS

- Your company has purchased a 3Mbps service contract to the Internet using the Ethernet connection between R2 and BB2. The provider for this Ethernet service does not limit your bandwidth to 3Mbps, but instead charges your company for any unicast traffic received by BB2 over this amount.
- Configure SW2 to ensure that R2 conforms to the 3Mbps rate.
- Do not use policing to accomplish this task.

2 Points

1.3 Traffic Filtering

- The company has experienced recent security issues with PCs in VLAN 17 trying to connect to each other using Windows file and print sharing. After attempting to get the IS department to disable file and print sharing on the PCs without success, you have been tasked with ensuring that PCs in VLAN 17 cannot talk directly with each other but still can communicate with other ports or interfaces in VLAN 17.
- Use the minimum configuration needed to complete this task.

2 Points

1.4 Keepalives

- Due to limitations with the Frame Relay service provider, the switches in the cloud do not inform each other when one of their local DLCIs changes status. Therefore, if one side of the Frame Relay connection goes down the other side's local Frame Relay switch will not be informed about the status change of the remote DLCI. This, in turn, will cause the DLCI on the remote end to remain active.
- To help protect against this problem, ensure that a Frame Relay failure can be detected by having R4 and R5 poll each other to ensure that the other side's Frame Relay interface is up and reachable every 15 seconds.

2 Points

2. IPv4

Note: Do not redistribute between IGPs for this lab.

2.1 OSPF

- Configure SW4 to match exactly the output below and nothing more:

```
Rack1SW4#show ip ospf database | include Net Link States \ (Area 34\)  
Summary Net Link States (Area 34)  
  
Rack1SW4#
```

2 Points

2.2 BGP Route Reflection

Routers are configured in BGP Autonomous Systems per the diagram. The following BGP peering sessions are pre-configured for you:

Device 1	Device 2
R2	BB2
R2	R3
R2	R4
R3	R1
R1	R4
R1	SW1
SW1	SW2
SW2	BB3
SW2	R5
R5	R4
R4	R6
R6	BB1

- R4 and R5 should be configured as route-reflectors in AS 100; These devices should treat each other as non-clients.
- R1 and R3 should be designated as route-reflectors in AS 200; These devices should treat each other as non-clients.

2 Points

2.3 BGP Origination

- Advertise VLANs 3, 17, and 33 into BGP on R1 and R3.
- Advertise VLANs 45, 46, and 58 into BGP on R4 and SW2.

2 Points

2.4 BGP Bestpath Selection

- Configure AS 200 so that all traffic destined for VLAN 3 uses the Ethernet segment between SW1 and SW2.
- All traffic destined for VLAN 33 should use the Frame Relay segment between R1 and R4.
- In the case that either of these links is down, traffic should be able to be rerouted out the other link.
- The Frame Relay circuit between R2 and R4 should not be used as transit to either of these destinations.

3 Points

2.5 BGP Filtering

- In order to avoid unnecessarily transiting additional devices in the path to AS 254, the BGP policy of AS 200 states that the only link that can be used to reach AS 254 is the Frame Relay circuit between R2 and R4.
- Under no circumstance should AS 100 be allowed to use its other connections to AS 200 as transit to AS 254, regardless if the Frame Relay circuit between R2 and R4 is down.
- This configuration should be done in AS 200.

2 Points

2.6 BGP Default Routing

- Due to the memory limitations of SW1, AS 100 has agreed to send AS 200 default information. However, since AS 200 still has additional connections to AS 100, it wants to make a better routing decision based on longer prefixes. Unfortunately, AS 100 has refused to maintain a complex filtering policy for AS 200. Therefore, they have decided to send AS 200 a full view along with a default out each BGP connection.
- Configure AS 100 to reflect this policy.

2 Points

2.7 BGP Default Routing

- Since SW1 does not have the memory capacity to take a full view of the BGP table, AS 200's BGP policy dictates that the only prefix it should take from AS 100 is the default.
- Additionally, ensure that SW1 is the most preferred exit point out of AS 200 for a prefix that no other device in AS 100 has a longer match for.
- Configure SW1 to reflect this policy.

2 Points

2.8 BGP Bestpath Selection

- Since AS 100 is already using a large portion of the bandwidth on the Frame Relay circuit between R2 and R4, AS 200 does not want to send traffic for a large amount of prefixes out this link.
- Configure AS 200 so that it will only send traffic out this link that is destined for AS 100 and its directly connected customers.
- Configure this filtering in such a way that it can account for an arbitrary amount of new customers that may be connected to AS 100 in the future.
- This link should still be able to be used to send traffic out to AS 100 if there are no other longer matches throughout the BGP domain, but should only be preferred as a default exit point if SW1's connection to AS 100 is down.

2 Points

2.9 BGP Bestpath Selection

- Since R1 does not have the memory capacity to take a full view of the BGP table, AS 200's BGP policy dictates that the Frame Relay circuit between R1 and R4 should be used for all prefixes that the other BGP speaking devices do not have a longer match for.
- This exit point may be used as a default connection, but only if both the Ethernet connection between SW1 and SW2 and the Frame Relay circuit between R2 and R4 is down.

2 Points

2.10 BGP Aggregation

- To ensure that your upstream peers (AS 54 and AS 254) have full IP reachability to your network, configure your border routers to advertise an aggregate block of your internal address space to these neighbors.
- In order to prevent unnecessary forwarding within your network, configure these border routers so that no other devices within your network see this aggregate address block.

3 Points

3. IPv6

3.1 IPv6 Addressing

- Enable IPv6 routing on R1, R2, R3, R4, and R6.
- Use the address 2001:CC1E:X:1::Y/64 for R1's Ethernet interface.
- Use the address 2001:CC1E:X:3::Y/64 for R3's VLAN3 Ethernet interface.
- Use the addresses 2001:CC1E:X:46::Y/64 for the Ethernet segment between R4 and R6.
- Use the addresses 2001:CC1E:X:23::Y/64 for the Serial connection between R2 and R3.

2 Points

3.2 IPv6 over Frame Relay

- Enable IPv6 on the Frame Relay segment between R1, R2, and R4 using the addresses 2001:CC1E:X:124::Y/64.
- Use link-local addresses in the format FE80::Y on these devices.

2 Points

3.3 EIGRP

- Configure EIGRP for the Ethernet link between R4 and R6.
- Create and advertise into EIGRP three additional Loopback interfaces on R6 with the following IPv6 addresses:
 - 2001:205:90:31::1/48
 - 2001:220:20:3::1/64
 - 2001:222:22:2::1/80
- Ensure that R4 does not advertise to R1 or R2 any IPv6 prefixes that originated in EIGRP and have a prefix length longer than 64, even if other networks are added in the future.

2 Points

3.4 OSPFv3

- Configure OSPFv3 area 0 on the Frame Relay segment between R1, R2, and R4, the FastEthernet interfaces of R1, and R3, and the serial line between R2 and R3.

2 points

3.5 IPv6 Redistribution

- Redistribute between OSPFv3 and EIGRP on R4.
- Ensure full reachability through the IPv6 enabled network, without using any default routes.
- Do not use the `redistribute connected` command for this task.

2 Points

4. MPLS VPN

4.1 VRF

R6 is preconfigured for a VRF named VPNB on the Fa0/1 interface. R5 is preconfigured for a VRF named VPNA.

- Configure SW4 for a VRF named TEST, and add interface Fa0/6 to this VRF as a Layer 3 port, using the address 10.0.0.10/24.
- Configure an OSPF process for this VRF, using the process id of 129 and area 0.

3 Points

4.2 MPLS

- Configure the connections between R4 and R5 and between R4 and R6 for MPLS, using LDP as the label protocol.

2 Points

4.3 VPNv4

- Configure BGP for a VPNv4 peering between R5 and R6. Redistribute as needed on R5 and R6 for the VRFs. Verify that SW4 can ping the VPNA loopbacks on R5.

3 Points

5. Multicast

PIM dense-mode is pre-configured on the following interfaces:

Device	Interface
R1	Fa0/0
R1	S0/1
R3	S1/2
R2	S0/1
R2	Fa0/0

PIM sparse-mode should be used on the following interfaces:

Device	Interfaces
SW2	VL58
R5	Fa0/0, Fa0/1
R4	Fa0/0, Fa0/1
R6	Fa0/0

5.1 Multicast Distribution

- There is a Windows Media Server located on VLAN 17 that is streaming a video feed into your network. This feed is using the multicast group address 225.25.25.25 and the UDP port 31337. Users in VLAN 22 have been complaining that they are unable to receive traffic for this group. After looking into the problem further, it seems that R3 is having issues with sending multicast packets out the PPP link to R2, but can send unicast and broadcast packets. Since you have been unable to determine why this is happening, you have opened a case with TAC, however, hosts in VLAN 22 need access to this group immediately.
- Configure your network so that these hosts can receive traffic from this group.
- Do not enable PIM on any additional interfaces to accomplish this.

3 Points

5.2 Static RP

- Create Loopback1 on R4 and R5 using the IP address 150.X.0.255/32.
- Advertise these interfaces into OSPF area 0 on R4 and R5.
- Configure R6 and SW2 to use R4 150.X.0.255 as the RP for all multicast groups.
- R4 and R5 should exchange information on active multicast sources with each other.

3 Points

6. Security

6.1 Traffic Filtering

- Recent security monitoring of your network has indicated that various unauthorized devices have been attempting to telnet to R6 and gain access to the CLI. The only legitimate device in your network that should be allowed to telnet to R6 is the NMS located at 129.X.46.100.
- In order to detect these unauthorized attempts as they occur, configure R6 to deny and log all attempts to access it via telnet.
- Ensure that your NMS can still access R6 via telnet.

2 Points

6.2 Router Hardening

- Your security manager is worried about frequent hacker attacks trying to bruteforce the access password for R2.
- In order to minimize the effectiveness of bruteforce attacks, configure R2 to block any login attempts for 5 minutes after 10 unsuccessful attempts in a minute.
- Any unsuccessful attempts should be logged.
- Ensure that users connecting from your internal network 129.X.0.0/16 are not affected by this configuration

3 Points

7. Network Services

7.1 Logging

- After telnet logging had been configured on R6, it had been determined that there are too many devices attempting to access it to keep track of just by looking at the console output. In order to store and parse these log messages at a later date, the syslog service has been enabled on the NMS. NMS is located at 2001:CC1E:X:1::100, where X is your rack number.
- Configure R6 to send its logged access-list hits to this device.
- A log message should only be generated once 10 access-list hits have been accumulated.

2 Points

7.2 NTP

- Configure NTP on all of your devices throughout the network in order to accomplish this.
 - R3 and R6 should be the NTP servers.
 - R1, R2, and SW1 should get their time from R3.
 - R4, R5, and SW2 should get their time from R6.
 - All devices in BGP AS 100 are physically located in Chicago, IL (CST -6), while all devices in BGP AS 200 are physically located in Reno, NV (PST -8).
- Configure these devices to reflect the appropriate time zone and daylight savings time configuration.
- Configure SW3 and SW4 in such a way that they will display the exact time and date of the last restart using the `show version` command.

3 Points

7.3 DNS

- As your network has grown, it has become increasingly difficult to keep track of all the IP addresses of your network devices.
- In order to ease your device management and identification, configure R3 to provide hostname to IP address mappings for your network devices.
- Configure R1, R2, and SW1 to use R3 as the DNS server.
- Add three "A" records to the DNS database, namely RackXR1, RackXR2, and RackXSW1, pointing to the respective devices Loopback0 IP addresses.

2 Points

7.4 Default Gateway Redundancy

- Configure R4 and R5 to represent a virtual gateway on VLAN45 with the IP address of 129.X.45.6
- R5 is the least loaded of the two routers, so it must receive about 70% of clients' traffic.
- R4 should respond to ARP requests sent by clients on VLAN45.

3 Points

8. QoS

8.1 Frame Relay Traffic Shaping

- VoIP users on VLAN 46 and behind BB2 have been complaining about intermittent voice cutouts when making phone calls. After further investigation, you have determined that the utilization of the Frame Relay circuit between R2 and R4 is well within normal parameters. However, it seems that the VoIP traffic is getting delayed behind larger data packets. To partly resolve this issue, your design team has asked you to configure Frame Relay Traffic Shaping to minimize the amount of delay that this VoIP traffic must endure.
- The Frame Relay circuit between R2 and R4 has been provisioned at 512Kbps; ensure that neither of these devices sends traffic beyond this rate on this circuit.
- Additional VCs on R4 should equally share the remaining bandwidth of its T1 interface to the Frame Relay cloud.
- In order to allow VoIP traffic to be interleaved between larger data conversations, ensure that the maximum time it takes to transmit a packet across the Frame Relay network is 10ms.

3 Points

8.2 Priority Queueing

- Now that the Frame Relay network is configured to conform to its provisioned rate, configure your network so that all VoIP traffic (UDP 16384 – 32767) for VLAN 46 that traverses the Frame Relay circuit between R4 and R2 gets priority over data traffic.
- VoIP should be allocated a maximum of 192Kbps during periods of congestion on this link.

3 Points

8.3 Traffic Marking

- There is a server in VLAN58 that hosts HTTP and STMP applications.
- Configure SW3 to mark HTTP responses entering on the port connected to R5 using DSCP value of AF21.
- At the same time, STMP reply traffic should be marked with DSCP value of AF23.
- In addition to that, limit the cumulative rate of HTTP and SMTP traffic to 2Mbps.

3 Points

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	16
IPv6	8
MPLS VPN	0
Multicast	5
Security	13
Network Services	16
QoS	17

GOOD LUCK!

1. Layer 2 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

Some IGP settings and IPV4 addressing have been preconfigured for you.

2.1 RIPv2

- Configure RIPv2 on R3.
- Enable RIP on the Ethernet segment between R3 and BB2.
- Configure R3 to authenticate all RIP updates received on VLAN 32 with a keyed hash value based on the password CISCO.
- Configure RIPv2 on R4, R5, and SW2.
- Enable RIP between R4 & SW2 and between R5 & SW2.
- Enable RIP on the Serial link between R4 & R5 and on R5's VLAN5.
- 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.
- 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.
- Configure your network so that RIP converges 10 times as fast as the default settings.
- Ensure to maintain the default timer ratio.

5 Points

2.2 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

2.3 IGP Redistribution

- Configure SW3 and SW4 for OSPF per the diagram provided.
- 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.

2 Points

2.4 BGP Traffic Engineering

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.
- 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.

4 Points

2.5 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 than 150000 prefixes 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.

3 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

3.3 Tunneling

- Connect R4 and R5 to the IPv6 network using ISATAP tunneling.
- The tunnel should connect R4 and R5 to R3 using the prefix 2001:CC1E:X:345::/64.
- Create IPv6 Loopback interfaces in R4 and R5 with the IP addresses 2001:CC1E:X:Y::Y/64 and make sure the rest of the routing domain could reach them.
- Ensure connectivity to the OSPFv3 routing domain using the minimum amount of static routes.

3 Points

5. 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 are 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 zero "giaddr" field.
- Assume R3 is the only device allowed to relay DHCP messages.

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. Network 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.
- Do not use the `host` command under DHCP pool settings to accomplish this.

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 "InternetNetworkExpert.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

7.6 Router Redundancy

- Enable gateway auto-discovery based on ICMP messages for the hosts on VLAN 367
- R3, R6 and SW1 should advertise themselves as candidate default gateways.
- R3 should be the preferred gateway for the hosts on the segment.
- The advertisements should be sent between the 10 and 30 seconds intervals.

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

IEWB-RS-VOL2 Lab 14

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	15
IPv4	29
IPv6	6
MPLS VPN	0
Multicast	8
Security	8
Network Services	9
QoS	4

GOOD LUCK!

1. Layer 2 Technologies

1.1 Packet Sniffing

- Users in VLAN 1011 have been reporting slow network response time, however, you have not been able to track down the problem. In order to collect more information regarding the issue, a Mac OS X host running *tcpdump* utility has been connected to port Fa0/12 of SW1.
- Configure SW1 so that all traffic received in VLAN 1011 is redirected to this host.
- The same port is to be used by the host to send regular data packets. The IP address assigned to the host is in VLAN 5.

2 Points

1.2 Configuration Management

- In order to protect against network downtime, your operations team has implemented a new policy which dictates that the current running configuration of SW1 must be archived in flash before any changes are made. Therefore, this file can be used as a reference against any newer configurations that cause problems in the future.
- This backup of the configuration should be stored in the *archive* directory and use the name *backup.config*.
- In order to make this process as simple as possible, configure SW1 so that when your administrators run the command *backup*, it automatically runs this process for them.
- Additionally, in order to minimize downtime in the event of a software crash due to faulty configuration changes, configure SW1 to load the archived backup configuration created upon the next bootup.

2 Points

1.3 Traffic Filtering

- In an effort to increase security, the network administrator has requested that SW1's port Fa0/5 be configured to only accept traffic from the MAC address 0000.0c12.3456.
- Configure SW1 to reflect this request, but do not use the command `switchport port-security mac-address 0000.0c12.3456` to accomplish this.

2 Points

1.4 Spanning Tree

- Disable spanning tree for VLAN 1363 on SW3 and SW4.
- Ensure that this does not create a spanning tree loop with SW4, since SW3 and SW4 are connected using two interfaces in VLAN 1363.
- Use the minimal configuration needed on SW4's interface Fa0/20 to accomplish this task without using the `shutdown` command.

2 Points

1.5 Switch Features

- Configure so that all traffic from R1 or R3 to another device on VLAN 1363 passes through SW3.
- Traffic from R1 to R3 should also pass through SW3.
- Do not add any VLANs or configure any trunks to achieve this task.

3 Points

1.6 Hub-and-Spoke

- Without creating subinterfaces, configure a Frame Relay hub-and-spoke network between R1, R3, and R5 with R3 as the hub.
- Traffic from R1 destined for R5 should transit R3, 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

1.7 Point-to-Point Addressing

- Configure R4 and R5's Serial interfaces to use the IP addresses 167.X.45.4/32 and 167.X.45.5/32 respectively.
- Ensure that R4 and R5 can ping each other's Serial interfaces.
- The creation of additional logical interfaces is permitted for this task.
- Do not use static routing to accomplish this.
- Do not use the command `peer default ip address` for this task.

2 Points

2. IPv4

2.1 RIP

- Configure RIPv2 on R4 on the Ethernet segment connecting to BB2.
- RIP updates received on this interface should be authenticated with a secure hash value of the password CISCO.
- Recently, you have been getting complaints from users about reachability problems to prefixes learned from BB2. After consulting with the administrators of BB2, it appears that your RIP updates are getting periodically lost when sent over VLAN 42. Coincidentally, you seem to remember a recent issue with the Catalyst switch not forwarding certain multicast packets as it should. In order to see if this is in fact the problem, configure R4 so that RIP packets are sent as a broadcast instead of a multicast as they go out to BB2.

2 Points

2.2 OSPF

- One of your concerns about this migration is sub-optimal routing due to link speeds higher than 100Mbps in your network. In response to this, the other business unit has agreed that the layer 3 EtherChannel between SW1 and SW2 should be seen with a cost of 10.
- Configure your network to reflect this policy.
- Ensure that all other link costs are automatically updated accordingly.

2 Points

2.3 EIGRP

- Do not allow BB3 to intercept EIGRP updates coming from any of the EIGRP speaking devices on VLAN 1363.
- Do not allow BB3 to send any EIGRP messages.
- Do not use the `neighbor` command to accomplish this.
- The use of VACLs is not allowed for this task.

3 Points

2.4 EIGRP

- Do not allow EIGRP to use more than 384Kbps of the 1.536Mbps T1 link between R4 and R5.
- Enable EIGRP on the Frame Relay network between R1, R3, & R5.
- Authenticate the EIGRP adjacency between R1 & R3 with the MD5 hashed password CISCO13.
- Authenticate the EIGRP adjacency between R3 & R5 with the MD5 hashed password CISCO35.

3 Points

2.5 EIGRP

- Recently, you have noticed some inconsistencies in the EIGRP topology of various devices throughout the network. After looking into this issue further, you have discovered that R1 is low on memory and has not been computing DUAL correctly. Until the exact cause of this problem can be located, configure R1 so that it can only be used to reach networks which are directly connected to it.
- Additionally, to reduce the effect of SIA queries, set the time that R1 would wait in active state to the minimum.

2 Points

2.6 IGP Redistribution

- Redistribute between EIGRP and RIP on R4. Send a summary to BB2 for the major subnet of your topology and one that covers the loopback networks.
- Since R5 is the only place where EIGRP and OSPF meet, there is no reason for these domains to have specific reachability information about each other. Configure R5 to generate a default route into both the OSPF and EIGRP domains.

2 Points

2.7 BGP Peering

Basic BGP settings have been pre-configured in your rack as outlined below:

- In order to reduce the amount of iBGP peering sessions that need to be maintained within AS 100, R3 has been chosen as a central point of distribution for all iBGP learned routes. Your design team has notified you that additional devices will be added to your BGP network within the near future. These devices will be assigned the Loopback 0 addresses of 150.X.9.9 and 150.X.10.10. In order to ease in the integration of these and future devices into your BGP domain, the design team has suggested that you configure all iBGP peers of R3 (R1, R4, R5, & R6) in the peer group *iBGP*. Members of this group should all share the following attributes:

```
remote-as 100
route-reflector-clients
send-community
update-source Loopback0
```

- In order to prepare for the upcoming additions to R3's iBGP peers, configure these new devices as part of the peer group, however, do not allow R3 to attempt to initiate the BGP session.

3 Points

2.8 BGP Peering

- R4 has recently been acquired from AS 200, however, its upstream peer in AS 254 has not yet updated its BGP configuration. In the meantime, configure R4 to peer with BB2 in such a way that BB2 thinks R4 is in AS 200.
- This adjacency should be authenticated with an MD5 hash value of the password CISCO.

1 Point

2.9 AS-Path Manipulation

- Soon after implementing this quick fix, you received a call from an engineer from AS 200 who stated that they have lost reachability to AS 254. After working together on the problem, you and this engineer have realized that routes you are learning from AS 254 are getting prepended with AS 200 in the AS-Path, and are subsequently getting dropped when entering AS 200 downstream. Configure your network to resolve this problem.

2 Points

2.10 BGP Bestpath Selection

- Advertise VLANs 4 and 5 into BGP.
- Traffic from AS 54 destined to these prefixes should come in the Frame Relay link between R6 and BB1.
- Traffic for these prefixes should only come in from BB3 if the link between R6 and BB1 is down.

2 Points

2.11 BGP Summarization

- Advertise the Loopback0 networks of SW1 and SW2 into BGP.
- Routers outside of AS 65078 should only see one route to reach these two prefixes.

2 Points

2.12 AS-Path Manipulation

- Since SW1 and SW2's only connection to the rest of the network is through AS 100, administrators of SW1 and SW2 have decided not to apply for their own BGP AS number. Instead, AS 100 has assigned them the locally significant AS number of 65078. However, since this is not a valid public AS number, it cannot be leaked out onto the Internet.
- Configure your network so that AS 65078 is stripped out of the AS path when updates are sent to AS 100's upstream peers.

2 Points

2.13 BGP Route Injection

- Due to AS 65078's aggregation policy, AS 100 cannot implement a detailed ingress traffic engineering policy. Despite requests from your network team for AS 65078 to stop this aggregation, they have continued to do so. In response to this, your network team has had no choice but to manually re-inject the prefixes which AS 65078 has aggregated.
- Configure your network so that traffic for SW1's loopback enters the Frame Relay link between R6 and BB1.
- Additionally, all traffic for SW2's loopback should enter the Ethernet segment between R3 and BB3.
- Ensure that all other routers throughout your domain only have the aggregate block for this address space that AS 65078 has originated.

3 Points

3. IPv6

3.1 IPv6 Addressing

- Create two new Loopback interfaces on R6 with the IPv6 addresses 2001:150:X:26::6/64 and 2001:150:X:2E::6/64.
- Enable IPv6 routing in R6 and ensure that EIGRP is running on both new Loopback interface.

2 Point

3.2 IPv6 Tunneling

- Configure an IPv6 over IPv4 tunnel between R4 and R6.
- This tunnel should remain up if R4 loses the connection to either R3 or R5.
- Use the network 2001:167:X:46::/64 for this segment.
- Minimize tunneling overhead with your solution.

2 Points

3.3 EIGRP

- Configure the network 2001:167:X:4::/64 on R4's connection to VLAN 4.
- Enable EIGRP on VLAN 4 and the tunnel between R4 and R6.
- R6 should advertise just a single summary prefix encompassing both configured IPv6 subnets to R4.
- Additionally, R6 should advertise a default IPv6 route to R4.

2 Points

5. IP Multicast

IP multicast routing has been enabled in R3, R4 and R5. PIM sparse-mode is active on VLANs 4,5 and 1363 as well as on the Frame-Relay segments between R3 & R4 and R3 & R5. Additionally, PIM Sparse-Mode runs in the Serial link between R4 and R5.

5.1 RP Assignments

- Configure R4 as the RP for all multicast groups throughout your network.
- Recently, you read of a multicast network attack in which rogue hosts were injecting false Auto-RP messages into the PIM domain. Configure your network so that R4's RP assignment cannot be preempted by any Auto-RP learned information.

2 Points

5.2 MBONE Connectivity

- Your network design team has informed you that they would like to connect to the MBONE with a DVMRP tunnel over the Internet.
- The *mrouted* host where the tunnel will terminate has an IP address of 220.20.3.192.
- This host will be expecting the tunnel to be originated from R4 with a source address of 192.10.X.4.
- Configure R4 to reflect this request.
- Ensure that R3, R4, and R5 can use DVMRP derived information for RPF checks on multicast sources.

3 Points

5.3 DVMRP Interoperability

- Multicast sources on VLANs 4 and 5 will be delivering multicast feeds to hosts on the MBONE.
- Configure R4 to advertise a single route for these two networks over your DVMRP tunnel to the MBONE.

3 Points

6. Security

6.1 DoS Protection

- After further investigating the slow response time to your web server (167.X.4.119), it appears that the server is undergoing a TCP SYN DoS attack. You have reported this attack to your upstream provider for them to take the appropriate action. In the meantime, configure R4 to be a proxy for all TCP sessions initiated to this server.
- R4 should tear down any TCP session that has been inactive for more than 30 seconds.
- Additionally, R4 should start closing half-open TCP sessions after they have exceeded 1000.
- Once the amount of half-open sessions has dropped below 500, R4 should stop closing them.
- Configuration for this task should not use any commands that include the word "inspect".

2 Points

6.2 Attack Mitigation

- While researching recent security bulletins, you have discovered your ISA web server in VLAN 4 is vulnerable to an attack from packets with malformed IP options headers.
- Configure R6 to prevent this type of attack by dropping all packets it receives from BB1 containing IP options.

2 Points

6.3 Infrastructure Security

- Change the enable secret to CISCO for R4.
- Using AAA, create two new users on R4 for administration via telnet.
 - A user named OPERATOR authenticated using the password of CISCO. Allow this user to configure any HTTP server settings and view the configured HTTP server settings. These capabilities should not be applied to any other users on R4.
 - A user named ADMIN with the password CISCO who can perform any configuration.
- Console access should not be affected
- Do not use the hierarchical privilege levels model to accomplish this task.
- Do not use the root view to accomplish this task.

2 Points

6.4 Traffic Monitoring

- Configure R6 for Traffic Export with the following criteria:
- Use a profile name of BB1
- For incoming traffic from BB1, sample 2% of the traffic.
- For outgoing traffic to BB1, sample 5% of the traffic.
- Traffic should be exported out the Fa0/0 interface with a destination MAC address of 60.60.60

2 Points

7. Network Services

7.1 Configuration Management

- Recently, your operations team has suggested a new policy of backing up router configurations to your internal web server. They have requested that you create a menu system on R6 as a test deployment for level 1 engineers in the NOC to backup configurations.
- These engineers will login to R6 via telnet with the username NOC and the password CISCO.
- Once they login the following menu should appear:

```
Menu for Level 1 NOC users
```

- ```
1. View Current Configuration
2. Backup Current Configuration
3. Exit
```

```
Choose your selection:
```

- The internal web server's IP address is 167.X.5.115 and will be expecting the username NOC and password CISCO to be received via SSL at port 8080.
- R6's current configuration should be saved in the directory CONFIGS and have a filename of R6\_CONFIG.txt on the web server.
- Ensure that the users can view the entire running configuration when they choose the first selection.

**2 Points**

## 7.2 NAT on a Stick

- Port Fa0/14 of SW1 connects to one of your client sites. Your design team has allocated this customer the IP address 167.1.27.2/24. Users at this site are using the private IP address space 172.16.0.0/24. Since this address space is not routable throughout your network, your client's onsite administrator has requested that you configure Network Address Translation (NAT) on the border router to hide their address space. Unfortunately, the device you are using to connect to this client is a Catalyst switch, which does not support NAT. After further investigation, you have discovered that the client does have an extra router onsite. Unfortunately, this router (R2) only has one Ethernet interface. Despite this fact, your operations team has left you with the task of determining an appropriate solution.
- Configure R2 so that these hosts can access the network. You are allowed to configure R2 for policy routing to accomplish this task.

**3 Points**

## 7.3 ICMP Error Reporting

- Traffic accounting has indicated that hosts in VLAN 5 are sending traffic to destinations that R5 does not have a route to, and that R5 is constantly informing these hosts that it cannot reach the destination in question. To reduce processor load on R5, configure it so that it only generates one of these error messages every five seconds.

**2 Points**

## 7.4 Gateway Redundancy

- SW3's default gateway is set to 204.12.X.100.
- Configure R1 to proxy for this IP address using HSRP.
- If R1's interface S0/0 is down, R6 should proxy for this IP address.
- If R6's interface S0/0 is down and R1's interface S0/0 is down, R3 should proxy for this IP address.
- Do not use the `standby` command with the `track <interface>` option on R1.

**2 Points**



## 8. QoS

### 8.1 Priority Queueing

- One of your company's executives has been complaining about slow network response time. After your manager promised this executive that the problem would have been fixed by the last network upgrade, concerns are growing about the future of your IT department. In order to appease this executive and save the department you have decided to prioritize all of his traffic as it exits out to BB2.
- The executive's host resides on VLAN 4, and has an IP address of 167.X.4.204.
- Configure your network so that traffic for this host has absolute priority over all other traffic as it exits out towards BB2.
- Do not use legacy priority queueing to accomplish this.

**2 Points**

### 8.2 Congestion Management

- Recently, your customers in AS 54 have mentioned that access to your public web server is very slow. After further investigation, you have discovered that there is congestion on the Frame Relay link to BB1.
- Configure the network so that traffic from the web server is guaranteed 50% of the bandwidth of the Frame Relay circuit to BB1.
- The web server's IP address is 167.X.4.119.
- Do not use a `policy-map` to accomplish this.

**2 Points**



# IEWB-RS-VOL2 Lab 15

## 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   | 7           |
| IPv4              | 13          |
| IPv6              | 8           |
| MPLS Technologies | 3           |
| Multicast         | 10          |
| Security          | 10          |
| Network Services  | 22          |
| QoS               | 6           |

# GOOD LUCK!

## 1. Layer 2 Technologies

*VTP domain CISCO has been configured on SW1, SW2 and SW3 with all three switches being VTP servers. At the same time, VTP domain IE is used on SW4. VLANs and IP addressing have been pre-configured according to the diagram supplied with the lab. All trunks are using the dot1q encapsulation.*

### 1.1 Switch Features

- Frames sent into the layer 2 domain from R4's interface Fa0/1 should use Tag Protocol Identifier of 0x8100 and a VLAN ID of 54; frames sent from Fa0/0 should the same TPID but use a VLAN ID of 45.
- As these frames are received by the layer 2 domain, an additional metro tag of 254 should be added. These frames should be delivered to interfaces Fa0/1.45 and Fa0/1.54 on R5.
- Ensure that R5 is able to ping R4's interface address, using a packet size of 1600 bytes.
- Your NOC engineers have been noticing minor outages that seem to coincide with the security team updating ACLs on SW4. You have informed these engineers that the switch is temporarily blocking traffic through the port that the ACL is being updated on. Although this is a normal and desirable case, they have requested that this behavior be disabled.
- Configure SW4 to meet this requirement.
- Network monitoring has indicated that BB3 is generating an unusually large amount of broadcast traffic on the link to SW3.
- While the problem is investigated configure SW3 to only allow 750Kbps of broadcast traffic inbound from BB3.
- BB3 will be connecting using 10Mbps Ethernet/half duplex; hardcode SW3's interface Fa0/24 for these settings.

**4 Points**

### 1.2 PPPoE

- Configure R4 and R5 for PPPoE for the connection on the 130.x.54.0 network R5 should be the server, and R4 should be a PPPoE client. R5 should hand out addresses in the 130.x.54.0/24 subnet.
- Make sure that R4 and R5 are still able to form an OSPF adjacency over this connection.
- Do not use the `network` command to accomplish this.

**3 Points**

## 2. IPv4

*Some IGP settings and IP addressing have been preconfigured for you.*

### 2.1 IGP

- Configure the EIGRP domain so that R1 prefers R2 to get to VLAN 3.
- This configuration should be done on R1.
- Do not use an offset-list or prefix-list to accomplish this.
- As a security precaution, your corporate policy dictates that OSPF LSA advertisements should not be sent out interfaces that connect to stub networks.
- Do not use the `passive interface` command to accomplish this.
- Configure R3 to reflect this policy.
- In order to prevent false routing information from being injected into the OSPF domain authenticate the adjacency between R1 and SW1 with the MD5 hashed password CISCO.
- Do not use the `ip ospf authentication message-digest` command on either of these devices.
- No other adjacencies should be authenticated.

**5 Points**

### 2.2 IGP Redistribution

- In order to obtain full reachability, perform the following:
  - Redistribute EIGRP AS 10 into EIGRP AS 100 on R6.
  - Redistribute between OSPF and EIGRP on R1, R3, and R4.
- Ensure routing tables' stability with your solution. .

**4 Points**

## 2.3 BGP Outbound Route Filtering

- Network monitoring of R3 and R4 has indicated high CPU utilization which appears to be related to the BGP process. After looking into the problem further engineers in AS 200 have noticed that a full BGP table is being learned from AS 100 and then many of these prefixes are getting withdrawn due to AS 200's filtering policy. Although many prefixes are being filtered out the border routers of AS 200 must still process all these updates before they can be discarded. In response to this, AS 200 has requested that AS 100 maintain an outbound filtering policy for prefixes advertised to AS 200, however engineers in AS 100 have refused to do so due to the large administrative overhead. After heated negotiations, engineers of AS 100 and AS 200 have agreed to implement BGP Outbound Route Filtering (ORF).
- Configure ORF on the peering session between R1 and R4.
- R1 should send only the following prefixes to R4:
  - 28.119.16.0/24
  - 28.119.17.0/24
  
- Do not apply any filter on R1 to accomplish this.
- Configure ORF on the peering session between R2 and R3.
- R2 should send only the following prefixes to R3:
  - 112.0.0.0/8
  - 113.0.0.0/8
  - 114.0.0.0/8
  - 115.0.0.0/8
  - 116.0.0.0/8
  - 117.0.0.0/8
  - 118.0.0.0/8
  - 119.0.0.0/8
  
- Do not apply any filter on R2 to accomplish this.
- Use the minimum amount of lines necessary in the prefix-list on R3 to accomplish this.

**4 Points**

### 3. IPv6

#### 3.1 IPv6 Features

- Configure an IPv6 over IPv4 tunnel between R2 and R5 using the network 2001:130:X:25::/64.
- This tunnel should not be sourced off any physical interfaces of the respective routers.
- Hosts on VLAN 26 should only select R2 as a default gateway.

**3 Points**

#### 3.2 IPv6 OSPF

- Configure OSPFv3 on VLANs 6, 26 and the Loopbacks of R2 and R6.
- R2 should advertise VLAN 5 to R6.
- Static routing is allowed to accomplish this.

**3 Points**

#### 3.3 IPv6 Routing

- Configure one static route on R5 to gain reachability to all of the networks attached to R2 and R6.
- This route should be as specific as possible any overlap the minimum amount of address space necessary to gain reachability.

**2 Points**

### 4. MPLS VPN

#### 4.1 VRF

- Configure a tunnel between R1 and R2, using the loopback0 interfaces as the tunnel endpoints.
- Add the tunnel interface to a VRF named TESTVRF, and configure the tunnel interfaces for the addresses 130.X.124.Y/30, where X is your rack number and Y is the device number.
- Verify that R1 and R2 can ping the other end of the tunnel.

**3 Points**



## 5. IP Multicast

*Multicast routing has been enabled in R1, R2, R3, R6, SW1 and SW2. PIM sparse-mode has been enabled on the following interfaces:*

| Device | Interface |
|--------|-----------|
| R1     | Fa0/0     |
| R1     | S0/0      |
| R2     | Fa0/0     |
| R2     | S0/0.124  |
| R2     | S0/0.234  |
| R3     | S1/0      |
| R3     | Fa0/0     |
| R3     | Fa0/1     |
| R6     | Fa0/0     |
| R6     | Fa0/1     |
| SW1    | Fa0/14    |
| SW1    | VLAN 17   |
| SW2    | Fa0/14    |
| SW2    | VLAN 8    |

### 5.1 RP Assignment

- Configure R3 to advertise itself as a candidate bootstrap router throughout the PIM domain.
- Configure R1 and R2 as candidate RPs.
- R1 should service the multicast groups 224.0.0.0 – 231.255.255.255.
- R2 should service the multicast groups 232.0.0.0 – 239.255.255.255.
- Use the minimum amount of access-list entries on both R1 and R2 to accomplish this.

**3 Points**

## 5.2 Multicast Filtering

- Recent traffic monitoring has indicated that users in VLAN 8 have been abusing network bandwidth by subscribing to high traffic multicast feeds.
- To help reduce the load on the network configure SW2 so that users in VLAN 8 can only belong to three multicast groups at a time.
- Additionally ensure that these users cannot join groups for which R2 is the RP.

**2 Points**

## 5.3 IPv6 Multicast

- Configure IPv6 multicast and add IPv6 addresses on the following interfaces as specified below:
  - R6's Fa0/0 interface
  - R4's Ethernet connection to R5 – 2001:130:x:45::4/64
  - R5's Ethernet connection to R4 – 2001:130:x:45::5/64

**2 Points**

## 5.4 IPv6 Multicast RP

- Configure R6 as the RP for IPv6 Multicast.
- Do not configure R1, R2 or R3 for any IPv6 multicast commands.
- The output of `show ipv6 pim bsr election` on R4, R5, and R6 should show R6's VLAN6 interface address as the BSR for IPv6, with a priority of 192.
- You may add addresses, routes, and logical interfaces if needed.
- Other than the interfaces listed in section 5.3, no FastEthernet or Loopback interfaces should run IPv6 PIM on R4, R5, or R6.

**3 Points**

## 6. Security

### 6.1 Attack Mitigation

- Recently you have noticed very high utilization on numerous devices throughout your network. After further investigation you have determined that various hosts in VLAN 5 are infected with a SQL worm. In order to reduce the load on your network while your network administrators install the appropriate patches configure R5 to contain this traffic.
- Hosts infected with this worm are sending out 404 byte packets destined for TCP port 1433.
- Ensure that other traffic on this port is not affected by this filter.
- Do not use an access-list to accomplish this.

**3 Points**

### 6.2 Firewall Feature Set

- In order to prevent hosts from being infected in the future you have decided to implement CBAC on R5's connection to BB2. This way hosts from outside your network cannot initiate sessions into your internal network, which reduces the risk of virii and worms entering the network.
- Configure R5 to only allow traffic to come in the connection to BB2 if it has been originated from inside your network.
- For connectivity testing purposes ensure that R5 can ping BB2.

**3 Points**

### 6.3 CBAC Optimization

- Based on the statistics, the average amount of concurrent connections across R5 is 2000-2200.
- Configure R5 to optimize the performance of CBAC inspection according to this number.

**2 Points**

## 6.4 Control Plane Security

- Configure R6 to drop any packets destined to the router with IP source route option (either loose or strict).
- Do not use control plane protection to accomplish this and do not drop any other IP options..

**2 Points**

## 7. Network Services

### 7.1 RMON

- Recently you have been trying to justify to your management the need for additional bandwidth on R1's WAN connection. However your manager does not believe that the current circuit is being utilized as much as you say it is. In order to show him the amount of congestion the interface is undergoing, configure R1 to generate an SNMP trap whenever the output queue length (`ifEntry.21`) of its Serial0/0 interface exceeds 750 packets.
- This MIB value should be sampled every 60 seconds.
- When there are more than 750 packets in the output queue R1 should generate the message "*WARNING: Frame Relay Circuit Congested*".
- When the value falls back to 100, an event should be generated that reads "*NOTICE: Frame Relay Circuit Within Normal Utilization*".
- The server to send these SNMP traps to is 130.X.17.100.
- This server will be expecting the community string to be IETRAP.

**3 Points**

## 7.2 Banners

- In order to facilitate in verifying BGP route propagation you have decided to allow unauthenticated telnet access to R6 so users can view the BGP table.
- Configure R6 so that when users telnet in they are immediately put into privilege level 1 without having to enter a username or password.
- Once the command line is active the following banner should be displayed:

```

AS 100 Route View Server #####
Use this device to view the Internet routing #
table from the perspective of AS 100 #
#####
```

**3 Points**

## 7.3 Telnet Control

- After opening up access to R6 your security team has become concerned about hackers using R6 as a launching point for their telnet sessions.
- Configure R6 so that once users telnet into R6 they cannot telnet back out to another device.
- Do not use the `privilege exec` command to accomplish this.

**2 Points**

## 7.4 Gateway Redundancy

- Recently a failure of the category 5 Ethernet cable attached to R6's VLAN26 interface resulted in severe network downtime for the users in VLAN 26.
- In order to prevent this problem from occurring in the future your design team has mandated that both R2 and R6 should be able to play the role of the default gateway for VLAN 26 depending on which of them is available.
- Configure your network so that R6 is the preferred default gateway for this segment.
- In the case that R6 is unreachable R2 should take over as the default gateway on this segment.
- If R6 returns after a failure R2 should relinquish its role as the default gateway for the segment. However in order to ensure that the routing domain has properly re-converged R6 should not assume the role of the gateway until it has been up for at least five minutes.
- Do not use HSRP to accomplish this and use the IP address of 130.X.26.1 for the virtual gateway.

**2 Points**

## 7.5 Gateway Redundancy

- Even after implementing the previous configuration you have received a report of downtime from hosts on VLAN 26. Apparently the Frame Relay circuit between R6 and BB1 was down, but hosts were still sending their traffic to R6. To avoid this problem configure R6 to track the state of the Frame Relay circuit to BB1.
- Since LMI may remain active even if the PVC to BB1 is inactive your design team has recommended that R6 track reachability to the route 200.0.0.0/24.
- If this route is unreachable by R6 then R2 should become the active gateway for hosts on VLAN 26.

**2 Points**

## 7.6 Traffic Accounting

- R5 is the firewall between your network and the hosts behind BB2.
- Your security team is interested in how many hosts are trying to initiate sessions into your network.
- Configure R5 to keep track of these hosts attempting to violate the filtering policy.
- To prevent this table using up all of R5's memory ensure that a maximum of 100 entries can exist in the table at any given time.

**2 Points**

## 7.7 Broadcast Forwarding

- Configure R2 and R6 to forward DHCP requests from the hosts on the segment VLAN26 to a fictive DHCP server at the IP address 130.X.3.100.
- Both routers should insert Option 82 in the forwarded requests.

**2 Points**

## 7.8 DHCP Relay

- Configure a secondary IP subnet 10.X.26.0/24 on the interfaces of R2 and R6 using their numbers for the last octet.
- Configure both DHCP relays to try using their secondary IPs for the *giaddr* field in case if the DHCP server is not responding to the requests using the primary IP address as the *giaddr*.

**2 Points**

## 7.9 EEM

- Configure R4 to send an email warning message from [router@ine.com](mailto:router@ine.com) to the address [noc@ine.com](mailto:noc@ine.com) when the *receive\_throttle* parameter for Fa 0/0 interface is incremented by 3 packets last 60 seconds.
- Send a copy of the email to [admin@ine.com](mailto:admin@ine.com); the subject should contain the text: "R4 Fa0/0 warning – receive\_throttle incremented over 3.
- The body of the message should contain the output of the `show interface` command for the respective interface.
- Use the SMTP server's IP address 130.X.3.100.

**4 Points**

## 8. QoS

### 8.1 Legacy QoS Conversion

- Convert this legacy QoS configuration in R5 to the MQC, and apply to the VLAN 52 subinterface. Assume the subinterface rate to equal the interface rate.

```
interface FastEthernet0/0
 custom-queue-list 1
 !
 queue-list 1 protocol ip 1 tcp www
 queue-list 1 protocol ip 2 tcp ftp
 queue-list 1 protocol ip 2 tcp ftp-data
 queue-list 1 protocol ip 3 tcp telnet
 queue-list 1 default 4
 queue-list 1 queue 1 byte-count 5000 limit 30
 queue-list 1 queue 2 byte-count 3000
 queue-list 1 queue 3 byte-count 500
```

**3 Points**

### 8.2. Priority Queueing

- Host accessing an audio feed from VLAN 17 have been complaining about poor audio quality and dropouts. After further investigation it appears that this traffic is getting delayed behind larger data packets when R1 sends it out to the Frame Relay cloud.
- In order to resolve this problem configure R1 so that this audio traffic is always sent before any other data traffic out the Frame Relay link.
- The server's IP address is 130.X.17.139, and is sending the audio feed as unicast to UDP port 8940.
- Do not use a `policy-map` to accomplish this.

**3 Points**



# IEWB-RS-VOL2 Lab 16

## 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 | 14          |
| IPv4                 | 24          |
| IPv6                 | 9           |
| MPLS VPN             | 0           |
| Multicast            | 6           |
| Security             | 5           |
| Network Services     | 11          |
| QoS                  | 10          |

# GOOD LUCK!

## 1. Layer 2 Technologies

*Layer 2 settings have been preconfigured for you. Some settings might be missing and it's up to you to troubleshoot and fix the network. Use the diagram as your main reference.*

### 1.1 Pruning

- Some time ago a new switch was installed in your network that had a high configuration revision number and it erroneously overwrote your entire VTP domain.
- In order to protect against this type of misconfiguration in the future your new corporate policy dictates that all switches must run in VTP transparent mode. However since SW1, SW2, SW3, and SW4 are not advertising VLAN information to each other they cannot participate in VTP pruning. This has resulted in a large amount of unnecessary broadcast traffic being sent over your trunk links.
- In order to solve this problem manually configure your network to behave as though VTP pruning has been enabled.
- Trunk only the necessary VLANs on SW2's trunk to R6.

**3 Points**

### 1.2 Metro Ethernet

- SW1 and SW2 have been preconfigured to provide transparent layer 2 transit for VLAN 45 between SW3 and SW4 using the metro tags of 100 and 200.
- Configure interfaces Fa0/13 - 14 on SW3 and interfaces Fa0/16 - 17 on SW4 as access links that forward traffic for VLAN 45.
- SW3 and SW4 should see each other via CDP on these interfaces.

**3 Points**

### 1.3 Bridging over Frame Relay

- Users on VLAN 16 and VLAN 22 are running a legacy application that only supports broadcast transmission. In order to support this application your design team has decided to bridge these two segments together. Configure your network to that traffic between these two segments can be bridged.
- Ensure that the rest of the routing domain can still communicate with these segments.

**3 Points**

### 1.4 L2 Protocol Tunneling

- Provide a protection against multicast packet flooding for the metro Ethernet connection between SW3 and SW4.
- Provider-Edge interfaces should drop STP packets once their rate exceeds 100 per second.
- Ensure that your configuration does not affect CDP or VTP packets.

**3 Points**

### 1.5 MPLS

- After reading some articles on the benefits of label switching, your CTO has decided to start a pilot project of MPLS on the network.
- Configure MPLS to run on R4 and R5
- Specifically the Ethernet link between R4 and R5 should be utilized for this transport.
- Do not use the “implicit-null” label in your network.
- R4 and R5 must originate label bindings for their own loopback addresses

**2 Points**

## 2. IPv4

*The customer is currently running OSPF Area 0 throughout their network. Some changes are necessary.*

### 2.1 OSPF

- Configure OSPF area 3457 on VLAN 45 between R4 & R5 and on VLAN 47 between R4 & SW1.
- Configure OSPF area 3457 on the Frame Relay network between R3, R4, and R5.
- The Frame Relay circuit between R3 and R4 has a provisioned rate of 1024Kbps, while the circuit between R3 and R5 is only provisioned at 512Kbps. Ensure that R3 takes this into account when computing OSPF metrics on this segment.
- Advertise the Loopback 0 addresses of these devices into area 3457.

**3 Points**

### 2.2 OSPF

- Configure OSPF area 51 on the 192.10.X.0/24 subnet between R1, R2, R6, and BB2.
- Advertise the Loopback 0 addresses of R1, R2, and R6 into OSPF area 51.
- In order to reduce the amount of prefixes necessary in the IGP tables throughout the routing domain, configure your network so that OSPF enabled devices outside of area 51 only see one route to R1 and R2's Loopback 0 networks. This route should be as specific as possible and not unnecessarily overlap any address space.

**3 Points**

### 2.3 OSPF

- Configure OSPF area 38 on VLAN 38 between R3 and SW2.
- Advertise the Loopback 0 address of SW2 into OSPF area 38.
- OSPF area 3457 connects to public areas of your network infrastructure. Since services offered in OSPF area 38 are of a business confidential nature your corporate policy dictates that devices in area 3457 should not have access to the resources of area 38.
- Configure R3 to reflect this policy.

**3 Points**

### 2.4 OSPF

- You have noticed very high CPU utilization on R3. After further investigation it appears that many consecutive changes in the OSPF topology are causing R3 to constantly run its SPF algorithm over and over. In order to help deal with this issue until the topology changes are diagnosed configure R3 so that it waits 4 seconds after receiving a link state update packet before running SPF.
- Additionally configure R3 so that it waits at least 10 seconds between consecutively running the SPF algorithm.

**3 Points**

## 2.5 BGP Communities

- To ease in the identification and traffic engineering of prefixes learned from their upstream peer, AS 100 has implemented a clearly defined routing policy based on community values. This policy states that prefixes learned from AS 54 should be tagged with community values as follows:
  - Prefixes originated in AS 54 and learned from BB1 should be tagged with the community *54:1*
  - Prefixes originated in AS 54 and learned from BB3 should be tagged with the community *54:3*
  - Prefixes not originated in AS 54 and learned from BB1 should be tagged with the community *X:1*, where X is the originating autonomous system.
  - Prefixes not originated in AS 54 and learned from BB3 should be tagged with the community *X:3*, where X is the originating autonomous system.
  
- Configure R6 to reflect this policy.

**3 Points**

## 2.6 BGP Bestpath Selection

- Configure your network so that R6 prefers to use the PPPoFR link for prefixes in the community *54:1*.
- Configure your network so that R6 prefers to use the Ethernet link to BB3 for prefixes in the community *X:3*.
- Do not use local-preference to accomplish this.

**3 Points**

## 2.7 BGP Bestpath Selection

- Configure AS 200 so that all traffic destined for prefixes in the *54:1* community come in the Serial link between R1 and R3.
- Configure AS 200 so that all traffic destined for prefixes in the *X:3* community come in the Serial link between R2 and R3.
- Do not use MED to accomplish this.

**3 Points**

## 2.8 BGP Bestpath Selection

- Configure AS 300 so that all traffic destined for VLAN 5 comes in the PPP link between R3 and R1.
- Traffic should be rerouted to the other PPP link in the case that the first fails.
- Do not use AS-Path prepending to accomplish this.

**3 Points**

## 3. IPv6

### 3.1 IPv6 Tunneling

- Configure an IPv6IP tunnel between R5 and R6 to connect their IPv6 segments.
- These tunnels should be sourced from their respective Loopback0 networks.
- Use the addressing format 2001:CC1E:X:56::Y/64.

**3 Points**

### 3.2 IPv6 Traffic Engineering

- Configure the network so that R3 sends IPv6 traffic from R6 destined for R5 directly to R5.
- Traffic from R5 back to R6 should go via R3.
- Do not modify any OSPF cost values to accomplish this.
- You are allowed to use policy-based routing to accomplish this.

**3 Points**



### 3.3 EIGRPv6

- Enable EIGRPv6 on all interfaces running IPv6.
- R6 should advertise the minimum amount of EIGRPv6 routes to R5 necessary for it to reach all of R6's prefixes.
- R6 should not advertise any address space that it does not have a more specific route for.
- R5 should receive only a single optimal summary prefix from R6.

**3 Points**

## 4. MPLS VPN

*No scenarios in this section.*

## 5. Multicast

### 5.1 RP Assignment

- Multicast servers are located on VLANs 45 and 63 in your network.
- The servers in VLAN 45 are sending to groups in the range of 224.0.0.0/5.
- The servers in VLAN 63 are sending to groups in the range of 232.0.0.0/5, with the exception of the administratively scoped range.
- Configure R3 to assign R4 as the RP for the servers in VLAN 45 and R6 as the RP for the servers in VLAN 63.
- In the case that R4 is unreachable R5 should be the RP for the servers in VLAN 45.
- Groups in the administratively scoped multicast range should not be distributed throughout the multicast domain in either a sparse or dense fashion.

**3 Points**

### 5.2 RP Security

- Your network security team is concerned with your RP information leaking outside of your internal network. To prevent this configure R6 so that your RP information is not advertised out to devices in VLAN 63.

**3 Points**

## 6. Security

### 6.1 Source Verification

- Your security team has expressed concerns with the possibility of traffic sent from spoofed IP addresses being received inbound on R2's connection to BB2.
- In order to protect against this vulnerability configure R2 to drop any packets without a verifiable source IP address that are received from BB2.

**2 Points**

### 6.2 Traffic Filtering

- Your security team has informed you that a large amount of traffic coming in from R6's connection to BB1 is being sourced from RFC 1918 address space.
- Configure R6 to drop this traffic when it is received.

**3 Points**

## 7. Network Services

### 7.1 Authentication Failure Message

- Recently your security team has reported that someone is attempting a brute force attack on various devices throughout your network. Apparently this person seems to know that routers which display a "% Login invalid" message do not have AAA enabled and is specifically targeting these devices. In order to help discourage these attacks in the future the security team has requested that your border routers be configured to display a custom authentication failed message.
- Users who fail to authenticate should be given the message below:  

```
"Authentication Failed. Username or Password was Incorrect"
```
- As an additional measure to thwart these attacks on your network in the future, configure these devices to disconnect a session after one failed login attempt.

**3 Points**

## 7.2 Authentication Prompt

- The security team has also recommended that when users telnet into the border routers they should be presented with the username prompt of "Login Name: " and the password prompt of "Passcode: ".
- Configure the border routers to reflect this.

**2 Points**

## 7.3 Port Redirection

- Further monitoring of R6 has shown that most of the brute force attacks are going to the IP addresses of the interfaces connected to BB1 and BB3. In order to distract hackers and analyze their attack techniques your security team has installed a VMware honeypot terminal in VLAN 16 with a blank root password.
- Configure R6 so that all telnet and SSH requests sent to its outside interfaces are redirected to the honeypot.
- This machine's IP address is 192.10.X.112.

**3 Points**

## 7.4 Address Manipulation

- Configure a new Loopback interface on R4 using the 154.X.44.0/24 subnet.
- Configure R4 to automatically source all telnet sessions off this new Loopback interface.
- Without advertising this Loopback, ensure that users on R4 can successfully telnet to all devices in your network.

**3 Points**

## 8. QoS

### 8.1. Frame Relay Traffic Shaping

- The Frame Relay interfaces of R3, R4, and R5 all physically support a speed of T1, however the Frame Relay circuits are not provisioned in this way. The circuit between R3 and R5 is provisioned at 512Kbps while the circuit between R4 and R5 is provisioned at 1024Kbps.
- Configure FRTS so that all end points of the network conform to the provisioned rate.
- Both spokes should be allowed to burst up to their access-rate if they have accumulated credit.

**3 Points**

### 8.2 Application Filtering

- Administrators of your network have been having Quake 3 tournaments during lunch. Your management has expressed that this is not a problem as long as they're not playing Quake during normal business hours.
- Configure R5 so that these administrators can only play Quake 3 before business hours, during their lunch break, and after hours.
- Work starts at 9am, ends at 5pm, and the lunch hour is noon to 1pm.
- The Quake 3 server is located on VLAN 5 with the IP address of 154.X.5.100 and is sending Quake 3 traffic out using UDP port 27960.
- The administrators that are playing are on located on VLANs 47 and 3003.
- Do not apply an access-group to any interface to accomplish this.

**3 Points**

### 8.3 Prioritization

- The administrators on VLAN 3003 have been complaining that they are getting Owned while playing Quake due to high ping times. Since they are only playing during off hours you have decided to help them out and decrease the latency of their packets. Configure your network so that the Quake 3 traffic coming from the server is prioritized on its way to VLAN 3003.
- This traffic should be allotted as much bandwidth as necessary.

**2 Points**

## 8.4 Link Efficiency

- Due to slow speed of the Serial links connecting R1 and R2 with R3, you have been tasked to come with a solution that improves bandwidth usage.
- Use the Lempel-Ziv based algorithm that is more CPU intensive but requires less memory usage.

**2 Points**



# IEWB-RS-VOL2 Lab 17

## 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 | 3           |
| IPv4                 | 16          |
| IPv6                 | 0           |
| MPLS VPN             | 0           |
| Multicast            | 6           |
| Security             | 23          |
| Network Services     | 17          |
| QoS                  | 14          |

# GOOD LUCK!



## 1. Layer 2 Technologies

*Layer 2 settings have been configured to match the diagram supplied with the lab. Refer to the switches configuration for more information.*

### 1.1 Fault Tolerance

- The Serial link between R4 and R5 will be used as a backup of the Frame Relay circuit between them.
- Configure the network in such a way that this link is activated if the Frame Relay circuit between these devices goes down at any point throughout the provider cloud.

**3 Points**

## 2. IPv4

*Some IGP settings and IP addressing have been preconfigured for you.*

### 2.1 OSPF

- Configure OSPF area 137 on VLAN 137 between R1, R3, and SW1.
- R1 connects to VLAN 137 with a FastEthernet interface, while R3 connects to VLAN 137 with a regular Ethernet interface. Configure the network so that SW1 takes this factor into account when it is computing the OSPF cost to reach destinations through these neighbors.
- Configure OSPF area 23 on VLAN 23 between R2, R3, & SW2, and VLAN 8 on SW2.
- Advertise the Loopback 0 interfaces of SW1 and SW2 into the OSPF domain, but do not use the `network` statement under the OSPF process to accomplish this.
- Since SW2 has no choice to reach external parts of the network other than to go through R2 or R3 it does not need specific reachability information about these external prefixes. However since R2 and R3 have various connections to the rest of the internal network it is advisable for SW2 to have reachability information about the internal OSPF network.
- Additionally since R2 only has the single low speed Frame Relay circuit that connects to the rest of the network SW2 should send all traffic to R3 that is destined for prefixes outside of the OSPF domain.
- Configure your network to reflect this specification.

**5 Points**

## 2.2 IGP Redistribution

- Redistribute between OSPF and EIGRP on R5.
- In order to ensure optimal routing within the OSPF domain configure your network so that routes redistributed in from EIGRP have a cumulative metric throughout the OSPF domain.
- Redistribute between OSPF and RIP on SW1.
- To minimize the amount of prefixes in the routing table configure your network so that all devices except SW1 see only one route for the prefixes learned from BB3 via RIP.
- This route should be as specific as possible and not unnecessarily overlap any address space.

**3 Points**

## 2.3 BGP

*BGP has been pre-configured according to the following description:*

- *Autonomous systems assigned per the diagram.*
  - *Fully meshed iBGP peerings within AS 100.*
  - *Fully meshed iBGP peerings within AS 200.*
  - *eBGP peering session between SW1 and BB3.*
  - *eBGP peering sessions between R1 and R5, R2 and R5*
  - *eBGP peering session between R5 and BB2.*
- 
- Configure AS 100 so that it cannot be used as a transit AS for customers in AS 54 to reach AS 200, and vice versa.
  - Do not use access-list, prefix-list, or AS-Path access-list filtering to accomplish this.
  - In order to avoid transiting the already congested Frame Relay circuit between R2 and R5 configure AS 100 so that it sends all traffic destined to AS 254 out the Frame Relay connection between R1 and R5.
  - The above configuration should be done on R2. Do not use local-preference or AS\_PATH manipulation to accomplish this.
  - Tune BGP timers in R5 to run the BGP scanner three times more often than by default.
  - Ensure that R5 waits for no more that 12 seconds before it switches to the write mode and starts sending updates.

**4 Points**

## 2.4 Optimized Edge Routing

- Configure AS 100 for Optimized Edge Routing to automatically balance the use of the uplinks to AS 200.
- R3 should be the master controller while R1 and R2 should be the border routers. Enable OER event logging.
- Authenticate all OER communications using the key string of CISCO.
- Enable automatic prefix monitoring for throughput and delay. The address information should be summarized based on the BGP table contents.
- Run periodic learning procedure every 1 minute for the duration of 1 minute.
- A prefix should be declared OOP if the measured delay is above 200ms.
- An exit should be declared OOP if the utilization is above 60%.
- Interface utilization changes should be detected as quickly as possible.
- OER should automatically distribute the load among the two edge links, assuming that both R1 and R2 Frame-Relay links bandwidth is 64Kbps.
- Control the prefixes via BGP to accomplish this. Use the local preference value of 5000.
- Advertise R4's and R6's Loopack0 interfaces into BGP for testing.

**5 Points**

## 3. IPv6

*No scenarios in this section.*

## 4. MPLS VPN

*No scenarios in this section.*

## 5. IP Multicast

### 5.1 PIM

- Users on VLAN 46 want to join the multicast group 227.69.53.7 that is going to be streamed into your network from BB3. However since these are the only users in the network that you want to receive multicast feeds you do not want to enable multicast routing everywhere in your network.
- Configure your network to accommodate these users without affecting any other devices in the transit path.
- Do not use any RP assignments to accomplish this.

**2 Points**

### 5.2 Multicast Testing

- To ensure that this setup will work before the multicast stream is injected, configure your network so that R4 will respond to ICMP echo requests sent from SW1's VLAN73 interface to the group 227.69.53.7.

**2 Points**

### 5.3 Multicast Filtering

- Your design team does not want any multicast streams to be delivered to hosts on VLAN 46 other than the 227.69.53.7 group coming from behind BB3.
- Configure SW1 to reflect this policy.
- Do not use the `ip multicast boundary` or `ip multicast rate-limit` commands to accomplish this.

**2 Points**

## 6. Security

### 6.1 Traffic Filtering

- Recent traffic monitoring of your network has indicated that various hosts from behind BB1 are performing port scans on your network. Configure R6 so that these hosts are denied entry into your network. The IP addresses of these hosts are as follows:
  - 200.0.1.2
  - 200.0.3.2
  - 200.0.3.10
  - 200.0.1.18
  - 200.0.3.26
  - 200.0.1.10
  - 200.0.3.18
  - 200.0.1.26
- Use the minimum amount of lines necessary to complete this task.
- Do not deny traffic from any other hosts.

**2 Points**

### 6.2 Attack Mitigation

- Recently monitoring of your web server on VLAN 5 has shown an inordinate amount of half open TCP sessions, possibly indicating a DoS attack. In order to reduce the load on the server while the possibility of attack is investigated configure R5 to that TCP requests sent to this server are limited to a maximum of 500Kbps.
- Do not use the MQC syntax to accomplish this.

**2 Points**

### 6.3 Admission Control

- Hosts on VLAN44 of R4 are running Cisco Trust Agent. Configure R4 so that traffic from these hosts is only allowed into the network if they have authenticated to your RADIUS server using EAP over UDP.
- The RADIUS server's IP address is 173.X.137.252.
- The RADIUS server will be expecting the request to be sourced from R4's Loopback0 and use the password of CISCO.
- Ensure that if additional RADIUS servers are configured that they will automatically use the password of CISCO.

**3 Points**

### 6.4 Authentication

- Configure PPP on the Serial links between R1 & R3 and R2 & R3.
- R3 should challenge R1 and R2 to authenticate via CHAP.
- Use the minimum amount of `username` commands on R3 to accomplish this.

**3 Points**

## 6.5 Traffic Filtering

- Ports Fa0/10 and Fa0/11 of SW1 connect to your web and mail servers respectively. Since they are in the same VLAN, your security administrators are concerned about one server being compromised and an attack being launched on the other from inside your network.
- In order to prevent this configure SW1 so that these servers cannot pass traffic between each other.
- Port Access-Lists are not allowed to accomplish this.

**2 Points**

## 6.6 Traffic Filtering

- As an additional protective measure configure SW1 so that an attacker who has compromised your servers can not circumvent your security by sending frames to random unicast and multicast MAC addresses.

**2 Points**

## 6.7 Traffic Filtering

- Ports Fa0/22 and Fa0/23 on SW2 connect to the legacy shared portion of your network. Recently you have been getting complaints from users in VLAN 137 about slow network response time. After further investigation you have determined that too many users are connecting to the hubs attached to SW2. In order to help alleviate this congestion while additional connections are added to your switch block a new policy has been implemented which states that maximum of 5 hosts can be connected to either of these ports at the same time.
- Configure SW2 to reflect this policy.
- Traffic received from excess hosts should be dropped.
- In order to ensure that inactive hosts do not unnecessarily take up one of these spots ensure that their MAC addresses are flushed out of the CAM table if they have been inactive for over 5 minutes.

**3 Points**

## 6.8 Trusts

- The security administrator decided that R4 has trust issues for security. VLAN 4 needs to be protected.
- Web traffic from VLAN 4 should be allowed out either WAN Serial interface as well as VLAN46. FTP traffic destined for R5's Loopback should only be allowed to go out the WAN-Serial zone.
- Any web traffic that is destined to go out the Fa0/0 interface should be subject to a deeper inspection that will disallow URL fields greater than 222 characters.
- You do not need to modify any routing for this task.

**3 Points**

## 6.9 Pass-Through Trust

- Any IP traffic between the WAN-Serial Zone and the WAN-Ethernet Zone should be allowed unhindered.

**3 Points**

## 7. Network Services

### 7.1 Logging

- After applying this configuration the server administrator has reported that R4 is overwhelming the syslog server when debugging is turned on.
- Configure R4 to limit the number of syslog messages to 10 per second.

**2 Points**

### 7.2 Telnet Logging

- Your network security team has expressed interest in tracking login attempts going to SW1 from outside your network.
- Configure SW1 so that all attempts to login to the command line via telnet are logged locally, except those coming from your internal network.
- Ensure that these log message remain in the case that the switch crashes.

**3 Points**



### 7.3 Address Translation

- For the purposes of security you do not want devices beyond BB2 to have specific reachability information about your network. Configure your network so that BB2 only has access to your network when hosts from inside initiate the connection.
- When BB2 receives this traffic it should all appear to have originated from the 192.10.X.5 address.
- Do not allow traffic originated from outside of your internal network to use this translation.

**3 Points**

### 7.4 Address Translation

- After configuring the above address translation you have been receiving reports from various users that they are not receiving any e-mail. In addition to this it appears as though users on the Internet cannot get access to your company's web server. Apparently while configuring address translation you forgot to account for these servers which are located on VLAN 5. In response to this your server administrators have updated the A records for your web server and the MX records in your DNS to point to the public address 192.10.X.5.
- Configure R5 so that users from behind BB2 can access both of these services.
- Your web server located on VLAN 5 has the IP address 173.X.5.100.
- Users will be accessing this server via normal TCP port 80 as well as SSL.
- Your e-mail server located on VLAN 5 has the IP address 173.X.5.100.
- Users will be accessing this server via both SMTP and POP3.

**3 Points**

### 7.5 Secure Shell

- Configure R1 so that an outside user may access the device connected to R1's AUX port via SSH.
- The outside user should be able to SSH into the port 2002 and authenticate itself with the username CISCO and the password of CISCO.
- All SSH-related events should be logged locally.

**3 Points**

## 7.6 Traffic Export

- A Traffic Analysis Module has been placed on VLAN 5. It can be reached at 173.X.5.100 and MAC 0010.1731.5100
- Sample one out of every 10 packets from each of the three WAN interfaces (either direction) and forward them to the device for analysis.

**3 Points**

## 8. QoS

### 8.1 Application Auditing

- Recently your network administrators have been getting complaints from users in VLAN 5 about a general network slowdown. During the investigation into this issue one of your administrators seems to recall overhearing some users in the lunch room talking about downloading mp3 files at work. You now suspect that peer-to-peer file sharing programs are the cause of this network slowdown. Since you are not sure that this is the case you still need to collect more information regarding the issue.
- Configure R5 to detect whether this peer-to-peer file sharing traffic is transiting the Ethernet segment of VLAN 5 and the Ethernet segment connecting to BB2.

**2 Points**

### 8.2 Application Filtering

- After applying the previous configuration you have in fact confirmed that an exorbitant amount of your bandwidth is being consumed by peer-to-peer file sharing applications. Also, it seems that users are running a wide range of these applications including Morpheus, LimeWire, and KaZaA version 2.
- In order to alleviate the congestion caused by these applications configure R5 to neither accept nor send traffic for these applications on the Ethernet segment to VLANs 5 and BB2.

**3 Points**

### 8.3 Traffic Policing

- You agreed to provide transit services between AS 54 and AS 254.
- In order to optimize network usage, you decided to implement bandwidth oversubscription in your network.
- The traffic contract with AS254 specifies the CIR of 512Kbps and the EIR of 384Kbps
- Implement traffic policing inbound on R5's connection to BB2 using dual-rate MQC policer.
- Mark the exceeding traffic with the DSCP value of CS0 and drop packets violating the traffic contract.
- The measurement interval in the contract is set to 30ms.

**3 Points**

### 8.4 Congestion Management

- The customer in AS 254 marks VoIP packets with DSCP value of EF.
- Configure R5's physical Frame-Relay interface so that the VoIP packets are given priority treatment but limited to 256Kbps.
- Do not allow the customer to use more than 512Kbps of the interface bandwidth.
- All other traffic should be scheduled using flow-based fair queuing and have DE-bit set in Frame-Relay frames.
- Make sure that your configuration only affects the DLCIs connecting R5 to R1
- Implement your solution using MQC syntax only.

**3 Points**

### 8.5 Adaptive Shaping

- R5's Frame-Relay link physical rate is 1024Kbps and the Frame-Relay SP agreed to provide the peak rate of 256Kbp to R4.
- At the same time, the SP only guarantees to deliver 128Kbps to every remote site.
- Implement the above traffic contract and ensure you provide 64Kbps to the VoIP bearer traffic, distinguished by DSCP EF marking.
- Ensure the shaper slows down to the CIR once voice traffic is detected and enables traffic fragmentation at the same time.
- Assume the physical port speed of R4 to be 512Kbps.

**3 Points**



# IEWB-RS-VOL2 Lab 18

## Difficulty Rating (10 highest): 7

### 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                 | 21          |
| IPv6                 | 3           |
| MPLS VPN             | 0           |
| Multicast            | 7           |
| Security             | 9           |
| Network Services     | 20          |
| QoS                  | 19          |

# GOOD LUCK!

## 1. Layer 2 Technologies

*Note: All Layer2 information has been pre-configured in this lab. SW3 and SW4 only require IPv4 reachability to each other.*

## 2. IPv4

### 2.1 EIGRP

- Enable EIGRP 10 on R6's Frame Relay connection to BB1.
- Authenticate the adjacency between these devices with an MD5 hash value that represents the password CISCO.
- Use key 1 for this authentication.
- Redistribute between EIGRP processes.

**3 Points**

### 2.2 EIGRP

- Configure SW2 so that hosts in VLAN 8 use VLAN 18 to reach all hosts with an even number in the third octet, while VLAN 58 is used to reach all hosts with an odd number in the third octet.
- Do not use the `distance` command to accomplish this.
- Ensure that traffic is rerouted within 5 seconds if SW2 loses connectivity to either R1 or R5.

**3 Points**

### 2.3 EIGRP

- Network monitoring has reported congestion on the Frame Relay circuits between R1, R3, and R5. After further investigation it appears that constant changes in the routing topology are causing EIGRP to consume half of the link bandwidth on the Frame Relay circuits.
- In order to help deal with this problem until the cause of the topology changes is tracked down configure your network so that EIGRP cannot use more than 10% of the bandwidth on these Frame Relay circuits.

**2 Points**

## 2.4 EIGRP

- Engineers in your network operations center have recently noticed that the %DUAL-3-SIA message has been periodically appearing in your syslog server logs. After further investigation you have determined that the constant changes in the EIGRP topology have been overwhelming R3's CPU, which in turn is delaying its replies to EIGRP query messages.
- In order to help manage this problem while the source of the topology changes is found configure routers in the EIGRP domain to wait up to 5 minutes for a response to an EIGRP query message.

**2 Points**

## 2.5 On-Demand Routing

- R4's only connection to the rest of the routing domain is through R5. Therefore it does not need specific reachability information about the rest of your network.
- Configure R5 so that it can learn about R4's stub networks via CDP.
- Ensure that hosts on VLANs 4 and 44 have connectivity to the rest of your network.

**2 Points**

## 2.6 BGP Peering

- After attempting in vain to establish the BGP peering session between R5 and BB2 you have called AS 254 to see what the problem is. After hours of escalation you have come to realize that the administrators of BB2 mistakenly configured your remote-as number as 200, and have failed to tell you that their BGP peering sessions require MD5 authentication. Luckily they have told you that the password for authentication is CISCO. However their remote-as configuration statement cannot be changed until the next maintenance window which is not scheduled for another few months.
- Configure R5 to peer with BB2 and support their configuration in the meantime.

**3 Points**



## 2.7 NLRI Advertisement

- To ensure that your upstream peers have full IP reachability to your internal network advertise all of your IGP learned networks into BGP.
- Do not use the `network` statement under BGP to accomplish this.

**3 Points**

## 2.8 BGP Reachability

- In order to reduce the memory utilization throughout your network your design team has opted not to run BGP on any device besides R5 and R6.
- Configure the network in such a way that these routers still have reachability to all BGP learned prefixes, but do not need to carry a full view of the Internet routing table.
- Ensure that this configuration does not withdraw any previously learned IGP information.

**3 Points**

## 3. IPv6

### 3.1 NAT-PT

- The network administrator has requested that R3 provide communication so that a host on VLAN3 running only IPv6 can communicate with one of your servers running only IPv4.
- The IPv6 host's address is 2001:CC1E:X:3::100.
- The IPv4 server's address is 156.X.8.100.
- The IPv6 host should see the IPv4 server as 2001:CC1E:FFFF::100.
- The IPv4 server should see the IPv6 host as 156.X.8.50.
- Configure R3 to reflect this request.

**3 Points**

## 4. MPLS VPN

*No scenarios in this section.*

## 5. IP Multicast

### 5.1 RP Assignments

- Configure R1 and R5 as candidate RPs for your multicast network via Auto-RP.
- Configure SW2 as the mapping agent for these RPs using its most reliable interface IP address.
- R1 should service the multicast groups 224.0.0.0 – 231.255.255.255.
- R5 should service the multicast groups 232.0.0.0 – 239.255.255.255.

**3 Points**

### 5.2 Multicast Testing

- There will be a multicast media server installed in VLAN 8 in the near future. In order to facilitate in testing your multicast routing before this server is installed, configure R3's interface Fa0/0 to join multicast groups 224.24.24.24 and 232.32.32.32.
- Ensure that R3 responds to ICMP echo requests sent from SW2's interface VLAN 8 destined for these two groups.

**2 Points**

### 5.3 Multicast Filtering

- After implementing the above configuration you have been getting complaints from users on VLAN 3 trying to access the multicast feed originated by the server in VLAN 8. After further investigation, you have determined that a device inside of AS 54 is mistakenly being used as the RP for this group. In order to prevent this problem from occurring in the future, configure your network so that the Auto-RP announce and discovery messages cannot be sent to or received from BB3.
- Do not use access-groups or `ip multicast rate-limit` command to accomplish this.

**2 Points**

## 6. Security

### 6.1 SSH

- Your security team has informed you that they are concerned about clear text telnet traffic being used to manage your Catalyst switches.
- Configure SW1 and SW2 so that they can be access remotely in a secure manner.
- The domain name used to generate RSA keys on SW1 and SW2 should be INE.com.
- For maximum security configure SW1 and SW2 with a key length of 2048 bits.
- Ensure that SW1 and SW2 can no longer be accessed via regular text telnet.

**3 Points**

### 6.2 Traffic Filtering

- Your security team has asked you to implement a filtering policy for hosts located on VLANs 4 and 44. Configure R4 to conform to this policy as follows:
  - Hosts in VLANs 4 and 44 should be able to initiate VoIP calls to any destination using the H.323 codec.
  - Hosts in VLANs 4 and 44 should be able to browse the web at ports 80, 443, and 8080.
  - The FTP server located at 156.X.4.40 should be allowed to accept active FTP sessions.
  - Traffic between VLANs 4 and 44 should be unfiltered.
- All other traffic from these segments should be dropped.

**3 Points**

### 6.3 PPP Authentication

- PPP is configured on the link between R4 and R5.
- R5 should authenticate R4 across this link, but R4 should not authenticate R5.
- Configure R5 to request CHAP authentication
- If CHAP authentication is rejected by R4, R5 should offer PAP authentication.
- Configure R4 to refuse CHAP authentication offered during the LCP negotiation.
- R4 should send the username of ROUTER4 and the password of CISCO for PAP authentication.

**3 Points**

## 7. Network Services

### 7.1 Syslog

- You have been tasked with configuring R2 to log all critical and below messages to a syslog server at IP address 156.X.8.100.
- In order to organize these messages the syslog server will be expecting R2 to use the facility local2.
- You suspect that someone may be tampering with R2's syslog messages on the syslog server itself. You believe that certain messages relating to configuration changes on R2 are being deleted by a NOC engineer in an attempt to circumvent your change control policy.
- Configure R2 to send its syslog messages in such a way that you can determine if any of R2's syslog messages have been deleted from the server.

**2 Points**

### 7.2 Logging

- After reviewing your syslog logs it seems that someone is in fact deleting messages from the server. In order to determine what type of messages are being deleted configure R2 to track the number and type of log messages being generated and store this information locally.

**2 Points**

### 7.3 DNS

- Recently your internal DNS server failed and your network administrators have asked you to configure R6 as a DNS server while your normal server undergoes repair.
- Configure R6 in such a way that when you issue the command `ping host` from any of your devices, where *host* is the hostname of any of your routers 1 through 6 or Switches 1 through 4, R6 receives requests for `host.ine.com` and resolves.

**3 Points**

### 7.4 Traceroute

- Recently administrators in your NOC have been complaining that it is too hard to decode the output from a traceroute going through your network. Apparently every time they traceroute they have to look at the IP addressing table to see which device has which IP address. They have requested that all devices in the network simply reply to a traceroute from their Loopback 0 interfaces. Although the other engineers on your team have told the NOC engineers that this is not possible, you know that it can be done. In order to show off your skills to your coworkers, configure R1 so that it always replies to a traceroute from its Loopback 0 interface.
- You are allowed to create an additional IP address to accomplish this.

**4 Points**

### 7.5 EEM

- If the transmit load value on R6's Frame-Relay interface goes over 1,000 Bytes per second, averaged over 4 seconds, EIGRP should be reconfigured to utilize "load" as part of the metrics.
- Use a chassis-id of RackX-R6 where 'X' is your rack number.

**4 Points**

## 7.6 Logging

- Other routers should detect the change to EIGRP and add “load” to their EIGRP processes as well

**2 Points**

## 7.4 Layer 2 Filtering

- Configure SW2 so that only PPPoE connections are allowed on VLAN8.
- Your solution should account for any future ports added to VLAN 8.

**3 Points**

## 8. QoS

### 8.1 Traffic Limiting

- Recently an Ethernet drop has been installed in your network as a new connection to the Internet. This link terminates at a public peering point, and is used to connect to both BB2 and BB3. BB2 will only allow R5 to send traffic across this link at a maximum of 2.5Mbps. BB3 will only allow R5 to send traffic into its network at a maximum rate of 3Mbps.
- Configure R5 to conform to these provisioned rates.

**3 Points**

### 8.2 Priority Queueing

- VoIP users connected to VLAN 4 have been complaining about poor voice quality when calling other users behind BB2.
- In order to help improve voice quality configure your network so that 64Kbps of bidirectional VoIP traffic is guaranteed to be dequeued first over the Serial link between R4 and R5.
- Additionally, to ensure that this VoIP traffic does not endure additional delay when sent out to BB2, configure R5 so that 64Kbps of this VoIP traffic is guaranteed to be dequeued first out the Ethernet link.

**3 Points**

### 8.3 Traffic Limiting

- As preventative maintenance against DoS attacks being launched from your network your security team has requested that you limit all ICMP traffic to a maximum of 16Kbps when implementing your QoS policy out to BB2 and BB3.
- Configure R5 to reflect this policy.

**3 Points**

### 8.4 DSCP Marking

- Lastly to try to fool BB2 and BB3 into providing your data traffic with expedited forwarding configure R5 so that all traffic sent out to both BB2 and BB3 is marked with a DSCP value of 101110.

**3 Points**

### 8.5 Policing

- Hosts attached to ports Fa0/10 and Fa0/11 on SW1 have been sending an inordinate amount of traffic into the network. Most of this traffic is specifically being set with DSCP values of EF and CS5.
- Configure SW1 to limit the reception of this traffic from these ports to 1Mbps.
- Traffic above this rate should be dropped.

**3 Points**

## 8.6 Per-Port Per-VLAN Policing

- Configure QoS marking in SW1 per the following requirements:
  - Set IP precedence to 3 for ICMP packets on VLAN52 and limit the input rate to 256Kbps.
  - Set IP precedence to 4 for TCP packets on VLAN53 and limit the input rate to 512Kbps.
  - Remark the exceeding TCP packets on VLAN53 with DSCP value of CS1.
- Packets enter the trunk link marked with DSCP value of CS0.
- Do not use the interface-level policing to accomplish this task.

**4 Points**



# IEWB-RS-VOL2 Lab 19

## 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 |
|----------------------|-------------|
| Layer 2 Technologies | 9           |
| IPv4                 | 12          |
| IPv6                 | 6           |
| MPLS VPN             | 0           |
| Multicast            | 9           |
| Security             | 5           |
| Network Services     | 26          |
| QoS                  | 12          |

# GOOD LUCK!

## 1. Layer 2 Technologies

### 1.1 Layer 2 Features

- Configure EtherChannel links between SW1 & SW2 and SW2 & SW3 using all available links. Do not use automatic negotiations to accomplish this.
- These links should use a 4-byte trunking encapsulation.
- Traffic leaving these links on SW2 should be load balanced based on the destination IPv4 address.
- Users in VLAN 127 have been reporting slow network response time, however your administrators have not been able to track down the problem. In order to collect more information your NOC engineers have requested that you redirect all traffic received in VLAN 127 on SW1 to a host running Ethereal in your network.
- This host is attached to port Fa0/10 of SW3.
- Use VLAN 10 for transporting this traffic.

**4 Points**

### 1.2 Frame-Relay Switching

- Configure R3 to provide a backup Frame-Relay link between R1 and R2.
- Use the Serial interfaces connecting R1 and R2 to R3 to accomplish this.
- Use the DLCI numbers 132 and 231 on the links connecting R1 to R3 and R2 to R3 respectively.
- Do not use the `frame-relay route` command to accomplish this.
- Do not use Inverse-ARP to establish IP connectivity.
- Refer to the diagram for the information on IP addressing.

**3 Points**

### 1.3 PPP

- Configure R4 to authenticate PPP session with R5 using the plain-text authentication.
- R4 should query the imaginary TACACS+ server at the IP address 149.X.4.100 authenticating and encrypting using the key of CISCO.
- Source TACACS+ packets off Loopback0 interfaces.
- In the case the TACACS+ server is not available, use the locally configured credentials.
- Use the name R5PPP along with the password of CISCO to authenticate R5.
- Do not use the global command `tacacs-server` host to accomplish this.

**2 Points**

## 2. IPv4

### 2.1 OSPF

- Configure OSPF area 0 on the Frame Relay connection between R3, R4, and R5.
- Configure your network so that R3 and R4 gain reachability to each other over the Frame Relay network through layer 3 routing instead of static layer 3 to layer 2 resolution.
- Advertise VLAN 44 into OSPF area 0.
- Recently a Windows host on VLAN 568 running OSPF injected false information into your routing domain and caused a traffic black hole. In response to this you have put a new policy in place which states that all adjacencies in OSPF area 568 must be authenticated with a secure hash value.
- In addition to the above, configure your network so that unauthorized devices cannot intercept OSPF hello packets as they are transiting VLAN 568.
- Administrators of your network have been noticing inconsistencies with the OSPF database when the Serial PPP link between R4 and R5 is being used. After further investigation they have determined that congestion on this link has been preventing LSAs from correctly propagating. In order to deal with this problem your design team has suggested that you increase the estimated time required to send a link-state update packet on this interface to 5 seconds.
- Additionally they have suggested that if an acknowledgement for an LSA sent across this interface is not received within 10 seconds, the LSA should be retransmitted.
- Configure the network to reflect this recommendation.

**4 Points**

## 2.2 BGP Features

- Advertise VLANs 4, 5, 7, 8, 77, 88, and 127 into the BGP domain.
- Advertise the Frame Relay network between R1, R2, and R3 into BGP.
- Advertise the Loopback 0 interfaces of R1, R2, and SW1 into BGP on R3.
- All of these prefixes should have an origin code of *incomplete* after being advertised into BGP.
- Since AS 300's only upstream peer is AS 200, it does not need specific forwarding information about the rest of the BGP domain.
- Configure your network so that AS 300 sees only a default route from R3, as well as prefixes originated by AS 200's directly connected customers.
- Configure AS 300 so that all traffic destined for VLAN 7 enters the Frame Relay circuit between R1 and R3 while all traffic destined for VLAN 77 enters the Frame Relay circuit between R2 and R3.
- R3 should load balance traffic destined for VLAN 127 amongst both Frame Relay connections to AS 300.

**5 Points**

## 2.3 BGP Aggregation

- In order to help reduce the size of the global BGP table AS 200 has decided to aggregate all networks learned from their customers.
- Configure R3 to originate an aggregate prefix that represents all of the VLANs that have been originated into BGP.
- When originating the aggregate address AS 200 should include an ordered set of the autonomous systems from which the subnets were originated.
- Furthermore since AS 300 will not accept a prefix that has its own AS in the path, the aggregate should only include AS 100 in the ordered set.
- Configure R3 to reflect this policy.

**3 Points**

### 3. IPv6

#### 3.1 IPv6 Addressing

- Create a new Loopback100 interface on R3 with the IPv6 address of 2001:220:20:3::3/64
- Configure IPv6 on the Frame Relay circuit between R1, R2, and R3 using the network 2001:149:X:123::/64.
- Configure IPv6 on VLAN 127 between R1 and R2 using the network 2001:149:X:127::/64.
- Hosts in VLAN 127 should use R1 as their default gateway.

**3 Points**

#### 3.2 RIPng

- Configure RIPng on all interfaces running IPv6.
- Hosts on VLAN 127 should prefer to use the Frame Relay PVC between R1 and R3 to reach the Loopback interface of R3.
- If this circuit is down they should be rerouted to R2's PVC to R3.

**3 Points**

## 4. MPLS VPN

No scenarios in this section.

## 5. Multicast

IP Multicast routing is enabled on R1, R2, R3, R4, R5, and SW1.

PIM sparse mode is configured on the following interfaces:

| Device | Interface |
|--------|-----------|
| R1     | Fa0/0     |
| R1     | S0/0      |
| R2     | Fa0/0     |
| R2     | S0/0      |
| R3     | S1/0      |
| R3     | S1/1      |
| R4     | Fa0/1     |
| R4     | S0/0/0    |
| R5     | Fa0/0     |
| R5     | S0/0/0    |
| SW1    | VL7       |
| SW1    | VL77      |
| SW1    | VL127     |

### 5.1 RP Assignment

- Configure SW1 to announce itself as a Rendezvous Point for the PIM domain.
- R3 should be responsible for group to RP mappings.

**3 Points**

### 5.2 Multicast Testing

- A multicast server located in VLAN 7 will be sending feeds to users in VLANs 4 and 5. In order to ensure that this configuration will be functional configure the network so that R4 and R5 respond to ICMP echo requests sent to the group address 224.1.1.1 sent from VLAN 7.
- Do not use tunneling or static RP assignments to accomplish this.

**3 Points**



### 5.3 Multicast Rate-Limiting

- Configure VLAN4 and VLAN5 interfaces so that any given multicast feed destined to 224.1.1.1 is limited to no more than 512Kbps
- All multicast traffic going out of the mentioned interfaces should be limited to 2Mbps

**3 Points**

## 6. Security

### 6.1 Traffic Policing

- Recent traffic monitoring in your network has indicated a suspiciously high amount of ICMP packets being received on R6's Frame Relay circuit to BB1. After further investigation it appears as though your network is undergoing a DoS attack.
- In order to reduce the impact of this attack on the rest of your internal network configure R6 to police all ICMP traffic received from BB1 to 8Kbps with the minimum possible burst.
- Do not use an access-list to accomplish this.

**2 Points**

### 6.2 Address Spoofing

- After reviewing your log files you have determined that the DoS attack came from hosts with spoofed private addresses.
- To help prevent this type of attack in the future configure your network so that traffic will not be accepted from BB1 if it has been originated from these hosts.

**3 Points**

## 7. Network Services

### 7.1 IOS Image Management

- Recently a security auditor downloaded all of your devices' configuration files via TFTP. Subsequently management has decided that TFTP is too insecure of a method to backup your devices' configurations. You have been tasked with setting up R3 to test out the new FTP server that will be used to backup devices' configurations.
- The FTP server's IP address is 149.X.5.100.
- The username for R3 to use is R3FTP and the password is CISCO.
- For security reasons you have setup the FTP server to only accept FTP sessions sourced from R3's Loopback 0 interface.
- Configure R3 to meet these requirements.

**3 Points**

### 7.2 Logging

- You have been tasked with setting up the edge routers (R3 & R6) with the following logging parameters:
  - The console should receive all severity 6 and below messages
  - Console messages should be rate-limited to 5 per second
  - Log severity 4 messages and below and store them in the routers' buffer
  - When users telnet in and execute the `terminal monitor` command they should receive all messages except "debugging"

**3 Points**

### 7.3 Line in Use Message

- Configure R5's VTY lines to display a "Line in Use" message of "Try back in 10 minutes" when an incoming telnet connection is attempted but all lines are full.

**2 Points**

## 7.4 Banner Messages

- Configure R5 so that when users telnet in the following banner is displayed where X is the incoming line number:

```
R5 is for use by authorized users only. You are on line
number: X.
```

- Do not enter the line number statically.

**2 Points**

## 7.5 HSRP

- Configure HSRP on R1 and R2 for hosts on VLAN 127 using the group name *HSRP*.
- These hosts will have their default gateway set to the IP address 149.X.127.254.
- R1 should be the preferred gateway unless it loses its connection to the Frame Relay cloud.
- Do not use object-tracking to accomplish this.

**3 Points**

## 7.6 DHCP Relay

- Configure R1 and R2 to forward DHCP requests from users on VLAN 127 to your DHCP server with the IP address 149.X.5.50.
- Ensure that only the active HSRP router forwards the DHCP request to this server.

**3 Points**

## 7.7 Network Stability

- Based on the router logs, R5's interface connected to VLAN 568 segments has been flapping too much recently.
- Configure so that the interface flaps do not affect the IGP stability, but do not block connectivity to VLAN 568 for more than 15 seconds.

**3 Points**

## 7.8 Convergence

- Configure R4 to detect the loss of physical connectivity to the switch on VLAN44 interface as soon as possible.
- Do not use the Ethernet keepalive mechanism (Loopback Frames) to accomplish that.

**3 Points**

## 7.9 BGP Policy Routing

- Configure R6 to advertise a prefix 150.X.66.0/24 into BGP.
- This prefix should only be advertised into BGP during the weekdays (Mon-Fri) from 9AM to 18PM.
- Do not use EEM to accomplish this.
- You are allowed to use one static route to accomplish this task.

**4 Points**

## 8. QoS

### 8.1 Frame Relay Traffic Shaping

- Recently you have been noticing drops on R3's Frame Relay PVC which connects to R2. Apparently your level-1 administrators failed to take into account the difference in port speeds between R2's 64Kbps interface and R3's T1 interface when configuring this circuit.
- In order to help alleviate congestion configure Frame Relay Traffic Shaping on R3 to reduce its average output rate on the circuit.
- R3 should attempt to average on output rate of 64Kbps on this circuit.
- In the case that R3 has accumulated credit it should be allowed to send a maximum of 12Kb of data in a single interval.
- Use the default Tc for this circuit.
- Do not MQC-based syntax to accomplish this task.

**3 Points**

## 8.2 Frame Relay Traffic Shaping

- Further monitoring of R3's Frame Relay circuit to R2 has indicated that the issue has been resolved. However now you have been getting complaints from users on VLAN 127 about horrible network response time. The complaints seem to have been coming from users on VLAN 127 that are using R1 as their default gateway. After speaking with the rest of your network team, it seems that no other recent configuration changes have been made regarding this circuit.
- Configure your network to resolve this problem.

**3 Points**

## 8.3 RSVP

- The connection between R4 and R5 is to be used for VoIP traffic, along with RSVP used for admission control and bandwidth reservation.
- The PVC between R4 and R5 has provisioned bandwidth of 512Kbps.
- Enforce this rate on both sides using the shortest Tc, and ensure that the use of RSVP for bandwidth admission is allowed on both sides.
- Allocate 50% of the PVC bandwidth to RSVP flows, and allow maximum of 64Kbps to every flow.
- RSVP should be supported at PVC level only, and the maximum number of reservable flows should be 4 out of 128 dynamic flows.

**3 Points**

## 8.4 RSVP Testing

- Simulate an RSVP flow between the Loopback0 interfaces of R4 and R5.
- The router that makes reservation should be R5 and the bandwidth requested should be 64Kbps.
- Request the controlled-load type of service from the network.
- Simulate the flow between UDP source and destination ports 16384.

**3 Points**



## IEWB-RS-VOL2 Lab 20

### 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 | 12          |
| IPv4                 | 27          |
| IPv6                 | 6           |
| MPLS VPN             | 0           |
| Multicast            | 5           |
| Security             | 11          |
| Network Services     | 11          |
| QoS                  | 6           |

# GOOD LUCK!



## 1. Layer 2 Technologies

### 1.1 Spanning-Tree Protocol

- Recently engineers in your network operations center have informed you that your switches are experiencing very high CPU utilization. After further investigation you have determined that too many resources are being dedicating to running individual instances of spanning-tree protocol on a per VLAN basis. To help reduce CPU utilization run three instances of spanning-tree protocol to service all VLANs assigned throughout your network.
- Configure your network so that only VLANs 1, 5, 12, and 107 are mapped to the first instance of STP.
- VLANs 27, 34, and 58 should be mapped to the second instance of STP.
- VLANs 46, 89, and 363 should be mapped to the last instance of STP.
- The name of this spanning-tree domain should be IESTP, and use a revision number of 10.

**3 Points**

### 1.2 Spanning-Tree Protocol

- Configure SW4 as the root bridge for STP instances 0 and 2.
- Configure SW1 as the root bridge for STP instances 1 and 3.
- VLAN 27 traffic from SW1 to SW2 should be sent over the 802.1q trunk link between SW1 and SW2; this configuration should be done on SW1.

**3 Points**

### 1.3 Spanning-Tree Protocol

- VLAN 363 traffic from SW2 to SW3 should use port Fa0/18 in the case if Fa 0/13 would go down.
- If port Fa0/18 is down it should use port Fa0/17.
- This configuration should be done on SW3.

**3 Points**

## 1.4 Bridging Over Frame Relay

- Recently a point-to-point T1 circuit has been provisioned between R1 and R3 in order to migrate R1 off of the Frame Relay network. Additionally, your provisioning department has put in an order for a new circuit to be turned up between R2 and R3 over the Frame Relay cloud. In preparation for this new setup in your network the design team has prematurely changed your IP addressing scheme to fit the new point-to-point circuit between R2 and R3. Unfortunately your change control policy dictates that an IP address change on any non-host device in the network must go through a long approval process. As a workaround in the meantime configure R1 to provide transit services for this segment.
- Ensure that R1 will route out the T1 circuit to reach this network once IGP connectivity has been established.
- Do not use the `bridge irb` command on R1 to accomplish this.

**3 Points**

## 2. IPv4

### 2.1 OSPF

- Configure OSPF area 0 on the Ethernet segment between R5 and SW2.
- Since there can not possibly be any other neighbors on this segment R5 and SW2 should not elect a DR or BDR.
- Ensure the OSPF dead timers are set to 40 seconds on this segment but do not use the `ip ospf dead-interval` or `ip ospf hello-interval` commands to accomplish this.
- Configure OSPF area 5 in VLAN 5.
- Advertise the VLAN 89 and SW2's Loopback0 interface into OSPF area 0.

**3 Points**

## 2.2 OSPF

- Configure OSPF type 1 authentication on the Frame Relay network.
- Use the password of CISCO for this authentication.
- Do not use the `area 345 authentication` command to accomplish this task.

**2 Points**

## 2.3 OSPF

- Configure OSPF type 2 authentication for all adjacencies in area 0.
- Use key number 1 and the password of CISCO.
- Do not use the `area 0 authentication message-digest` command to accomplish this task.

**2 Points**

## 2.4 OSPF

- One of the design engineers has recommend that when R3 and R4 boot up that they should not used as transit routers until they have had time to fully synchronize their OSPF databases.
- Configure R3 and R4 to advertise all OSPF routes with a maximum metric for the first 10 minutes after they have booted up.

**2 Points**

## 2.5 OSPF

- One of your design engineers has reported to you that the both the CPU utilization and the link utilization of routers connected to the Frame Relay cloud is spiking roughly every 30 minutes. After explaining to this engineer that this is OSPF's 'paranoid update', and is normal behavior, he has recommended to the rest of the network team that OSPF be replaced with static routes. Since you have attended Internetwork Expert's CCIE Routing & Switching Advanced Technologies Class you once again inform this engineer that there is a very simple solution to this problem.
- Configure your network to resolve this issue.

**2 Points**

## 2.6 OSPF

- Configure OSPF area 345 on the Ethernet link between R3 and R4.
- Traffic from SW2 to VLAN 34 should use the Frame Relay circuit between R4 and R5.
- This configuration should be performed on R5.
- Do not use the `cost` or `bandwidth` keywords to accomplish this.

**3 Points**

## 2.7 IGP Redistribution

- Redistribute between EIGRP, RIP, and OSPF on R3.
- Redistribute between RIP and OSPF on R4.
- R6 should use R3 to reach routes inside the EIGRP domain, and use R4 to reach routes inside the OSPF domain.
- Ensure R6 sees R3's and R4's Loopback0 interfaces with the prefix length of /24.

**3 Points**

## 2.8 BGP Bestpath Selection

*BGP peering sessions have been pre-configured for you. Refer to the diagram for BGP AS numbers information and discover the active BGP peering sessions using the show commands. Notice that you will need working IGP in order for BGP to work.*

- Even though AS 300 is directly connected to AS 54, the fastest path to reach it is out through AS 100's OC3 link. In order to follow this forwarding path, configure your network so that all traffic destined for prefixes learned from AS 54 traverses the Ethernet segment between R4 and R6.
- In the case that the Ethernet segment between R4 and R6 is unavailable, AS 300 should reroute to R6 by using Ethernet segment between R3 and R6.
- Do not alter the weight or local-preference values of these prefixes to accomplish this.

**3 Points**

## 2.9 BGP Filtering

- After failed negotiations between management groups AS 200 has now refused to provide transport for AS 300 to reach AS 254.
- Configure AS 200 to reflect this policy, but do not use any outbound filtering techniques or community-based filtering.
- Ensure that R1 still has reachability to AS 254.

**3 Points**

## 2.10 BGP Redistribution

- To ensure that non BGP speaking devices have full connectivity your design engineers have recommended that R3 and R4 redistribute their BGP learned prefixes into IGP. You have voiced your concerns about redistributing the full BGP table into IGP and have suggested instead that R3 & R4 inject a default route. After further negotiations with the design team, you have agreed to redistribute BGP into IGP, but only those prefixes which are less than four autonomous systems away.
- Configure R3 and R4 to reflect this policy.
- To help safe guard this redistribution policy, configure R3 and R4 to reset any BGP session that is sending more than 1000 prefixes.

**3 Points**

## 3. IPv6

### 3.1 IPv6 Addressing

- Enable IPv6 processing on R2 and R5.
- Configure IPv6 on VLAN 27 using the network 2002:8EXX:3502:0027::/64 where XX is your rack number.
- Configure IPv6 on VLAN 5 using the network 2002:8EXX:0505:0005::/64 where XX is your rack number.

**3 Points**

### 3.2 IPv6 Tunneling

- Hosts on VLANs 5 and 27 want to talk to each other via IPv6. Additionally your design team has notified you that hosts on these segments will soon be communicating with other IPv6 enabled hosts outside your network as well. However, your current demand for IPv6 does not dictate that the protocol should be enabled on every device throughout your transit network.
- Configure your network in such a way that hosts on VLANs 5 and 27 can communicate with each other, and so that they can communicate with an arbitrary number of IPv6 enabled segments that are reachable via the IPv4 network in the future.

**3 Points**

## 4. MPLS VPN

*There are no scenarios in this section*

## 5. Multicast

### 5.1 PIM

- Configure IP Multicast routing on R1, R3, and R4.
- Configure PIM sparse mode on the following interfaces:

| Device | Interface |
|--------|-----------|
| R1     | Fa0/0     |
| R1     | S0/1      |
| R3     | Fa0/0     |
| R3     | S1/2      |
| R4     | Fa0/0     |
| R4     | Fa0/1     |

- Configure R4 to announce its Loopback 0 interface as the RP for all multicast groups.
- Do not use the `ip pim autorp listener` command to accomplish this.

**3 Points**

### 5.2 Multicast Testing

- Configure R1's Ethernet interface to join multicast group 231.31.31.31.
- R3 and R4 should be able to successfully ping the multicast group address joined by R1.

**2 Points**

## 6. Security

### 6.1 Traffic Filtering

- Recent traffic monitoring of your network has indicated that various hosts from behind BB1 are performing port scans on your network. Configure R6 so that these hosts are denied entry into your network. The IP addresses of these hosts are as follows:
  - 51.3.0.1
  - 51.5.0.1
  - 51.7.0.1
  - 51.3.0.9
  - 51.5.0.9
  - 51.7.0.9
- Use the minimum amount of lines necessary to complete this task.
- Do not deny traffic from any other hosts.

**3 Points**

### 6.2 Reflexive Access-Lists

- The majority of these port scans were destined to hosts on VLAN 27. In order to protect hosts on this segment in the future your security team has asked you to implement a reflexive access-list on R2.
- Configure this access-list on R2 in such a way that hosts using TCP and UDP based applications on VLAN 27 can access the rest of the network.
- Ensure that hosts outside VLAN 27 can access your web server at 142.X.27.100, and that you can ping and telnet to SW1/SW4's SVI for management purposes.
- Ensure the traceroute command works when sourced off VLAN27 subnet.

**3 Points**



### 6.3 EAP

- The Serial link between R1 and R3 uses PPP encapsulation.
- Your company has decided to migrate away from Challenge Handshake Authentication Protocol for all PPP links and implement the newer Extensible Authentication Protocol. Management has requested for R1 and R3's previous CHAP configuration be converted over to EAP.
- R1 and R3's configuration related to CHAP is as follows:

```
R1:
interface Serial0/1
 encapsulation ppp
 ppp chap hostname ROUTER1
 ppp chap password 0 CISCO
```

```
R3:
username ROUTER1 password CISCO
!
interface Serial1/2
 encapsulation ppp
 ppp authentication chap
 ppp chap hostname ROUTER3
```

**3 Points**

### 6.4 DDoS Attack Tracking

- The server with the IP address 142.X.34.100 is under distributed DoS attack.
- Configure R3 to store the IP addresses of the hosts flooding the server.
- The number of entries collected should not exceed 100

**2 Points**

## 7. Network Services

### 7.1 SNMP

- Two new network management servers have been installed to manage R5. Configure R5 for the following SNMP parameters:
  - Contact: CCIE Lab R5
  - Location: San Jose, CA US
- The first network management server's IP address is 142.X.5.100 and the second network management server's IP address is 142.X.58.100.
- The network management servers are expecting the RO community string to be CISCORO and the RW community to be CISCORW.
- SNMP traps should be sent with the community CISCOTRAP.
- Log any other device that tries to poll R5 via SNMP.
- To maintain consistency in monitoring R5's interfaces ensure that the `ifIndex` values do not change across reboots.

**2 Points**

### 7.2 SNMP

- After the installation of the two new network management servers, you have noticed high CPU utilization related to the SNMP process on R5. After further investigation it seems that the NOC is polling for R5's routing table and ARP table via SNMP.
- Disable the ability of R5 to be polled via SNMP for its routing table (ip.21) and ARP table (ip.22).
- R5 should continue support for all other MIBs (iso).

**2 Points**

### 7.3 IOS Image Management

- During a maintenance window the previous night you noticed that R3 had to be reloaded three times to finally get it to recognize its flash memory. This in turn caused R3 to try and boot a default IOS image via TFTP. Since most of your companies networking infrastructure was purchased off eBay you are not able to RMA the flash module with Cisco. Until you can buy a new flash memory module off eBay configure R3 to boot a default IOS image from R4 in the event that it can not locate its own image in flash.
- Do not apply any configuration on R3 to accomplish this task.

**2 Points**

### 7.4 Local Authorization

- You have opened a case with TAC to help troubleshoot an issue relating to R4 crashing. TAC has requested access to R4 in order to help troubleshoot the problem. Allow TAC to telnet into R4 using username TAC and password CISCO.
- Since your corporate policy denies non-company personnel access to your networking infrastructure, you have decided to only give TAC limited access. When the TAC engineer telnets into R4 they should be placed into privilege level 0 and given access to the following commands:

```
show version
show processes cpu
show stack
show memory
```

**3 Points**

## 7.5 Telnet Filtering

- The TAC engineers will be telneting from the following IP addresses:
  - 45.194.169.115
  - 61.202.173.243
  - 41.234.41.250
- Without regards to overlapping additional IP addresses use the most efficient one line access-list to permit these three IP address to telnet into R4.

**2 Points**

## 8. QoS

### 8.1 IP TOS

- Prior to implementing a new QoS policy, you have been monitoring your network for any packets that have the TOS byte set. You have noticed that TCP packets sourced by the routers have the first two most significant bits of the TOS byte set in the IP header. At first you thought these were just BGP packets and were not really concerned, but after looking closer you noticed that these were actually telnet packets. Since marking telnet packets with the TOS of 0xC0 will conflict with your new QoS policy, you have decided to have all routers set the TOS for telnet packets to 0x0. Configure your network to reflect this policy.

**2 Points**

### 8.2 WRED

- Users on VLAN 27 have been complaining about slow access to the rest of the network. After further investigation you have determined that the output queue of R2's Serial interface is full, and traffic attempting to enter the queue is getting dropped.
- To help alleviate congestion configure R2 to selectively drop traffic on the Serial interface before the output queue becomes full.
- Traffic with a higher DSCP value should be less likely to be dropped than traffic with a lower value.

**2 Points**

### 8.3 Marking

- After implementing the new queueing strategy on R2 you have noticed slow response time to your web server located on VLAN 27. Apparently the web server service is not marking its TCP traffic with a DSCP value, and is therefore less preferred over other traffic.
- To decrease response time to the server configure R2 so that traffic from this server is least likely to be dropped as it is sent out to the Frame Relay cloud.
- The server's address is 142.X.27.100.

**2 Points**