



CCIE Routing and Switching v4.0 Configuration Practice Labs

Second Edition

Martin J. Duggan

Cisco Press

Copyright.....	1
About the Author.....	2
About the Technical Reviewer.....	2
Acknowledgments.....	3
Introduction.....	4
Practice Lab 1.....	9
Equipment List.....	9
Setting Up the Lab.....	10
Pre-lab Tasks.....	13
Practice Lab One.....	14
Section 1: LAN Switching and Frame Relay (28 Points).....	15
Section 2: IPv4 IGP Protocols (22 Points).....	18
Section 3: BGP (14 Points).....	21
Section 4: IPv6 (14 Points).....	22
Section 5: QoS (8 Points).....	24
Section 6: Security (6 Points).....	25
Section 7: Multicast (4 Points).....	25
IP Services (4 Points).....	25
“Ask the Proctor”.....	26
Section 1: LAN Switching and Frame Relay.....	26
Section 2: IPv4 IGP Protocols.....	28
Section 3: BGP.....	30
Section 4: IPv6.....	31
Section 5: QoS.....	33
Section 6: Security.....	34
Section 7: Multicast.....	34
Section 8: IP Services.....	34
Lab Debrief.....	36
Section 1: LAN Switching and Frame Relay (28 Points).....	36
Section 2: IPv4 IGP Protocols (22 Points).....	47
Section 3: BGP (14 Points).....	63
Section 4: IPv6 (14 Points).....	74
Section 5: QoS (8 Points).....	88
Section 6: Security (6 Points).....	94
Section 7: Multicast (4 Points).....	98
IP Services (4 Points).....	101
Lab WRAP-UP.....	104
Practice Lab 2.....	105
Equipment List.....	105
Setting Up the Lab.....	106
Pre-lab Tasks.....	110
Practice Lab Two.....	111
Section 1: LAN Switching and Frame-Relay (24 Points).....	112
Section 2: IPv4 IGP Protocols (28 Points).....	114
Section 3: BGP (15 Points).....	117
Section 4: IPv6 (12 Points).....	119
Section 5: QoS (6 Points).....	121
Section 6: Multicast (7 Points).....	121
Section 7: Security (7 Points).....	121
“Ask the Proctor”.....	122
Section 1: LAN Switching and Frame-Relay.....	122
Section 2: IPv4 IGP Protocols.....	123
Section 3: BGP.....	126
Section 4: IPv6.....	126
Section 5: QoS.....	126
Section 6: Multicast.....	127
Section 7: Security.....	127
Practice Lab Debrief.....	128
Section 1: LAN Switching and Frame-Relay (24 Points).....	128
Section 2: IPv4 IGP Protocols (28 Points).....	136
Section 3: BGP (15 Points).....	156
Section 4: IPv6 (12 Points).....	165
Section 5: QoS (6 Points).....	174
Section 6: Multicast (7 Points).....	176
Section 7: Security (7 Points).....	180
Lab WRAP-UP.....	184
Practice Lab 3—The VPN Lab.....	185
Equipment List.....	185

Setting Up the Lab.....	186
Pre-Lab Tasks.....	189
Practice Lab Three.....	191
Section 1: LAN Switching and Frame Relay (6 Points).....	192
Section 2: MPLS and OSPF (19 Points).....	194
Section 3: BGP (5 Points).....	197
Section 4: EIGRP and MP-BGP (9 Points).....	198
Section 5: OSPF and MP-BGP (9 Points).....	199
Section 6: MPLS (7 Points).....	200
Section 7: VPLS Simulation (10 Points).....	200
Section 8: Multicast (10 Points).....	200
Section 9: IPv6 (6 Points).....	201
Section 10: QoS (13 Points).....	201
Section 11: Security (13 Points).....	202
Practice Lab 3: "Ask the Proctor".....	202
Section 1: LAN Switching and Frame Relay.....	202
Section 2: MPLS and OSPF.....	203
Section 3: BGP.....	203
Section 4: EIGRP and MP-BGP.....	204
Section 5: OSPF and MP-BGP.....	204
Section 6: MPLS.....	205
Section 7: VPLS Simulation.....	205
Section 8: Multicast.....	206
Section 9: IPv6.....	206
Section 10: QoS.....	206
Section 11: Security.....	207
Practice Lab 3 Debrief.....	208
Section 1: LAN Switching and Frame Relay (6 Points).....	208
Section 2: MPLS and OSPF (19 Points).....	211
Section 3: BGP (5 Points).....	223
Section 4: EIGRP and MP-BGP (9 Points).....	225
Section 5: OSPF and MP-BGP (9 Points).....	230
Section 6: MPLS (7 Points).....	234
Section 7: VPLS Simulation (10 Points).....	240
Section 8: Multicast (10 Points).....	244
Section 9: IPv6 (6 Points).....	248
Section 10: QoS (13 Points).....	252
Section 11: Security (13 Points).....	254
Lab 3 Wrap-Up.....	262
Chapter 4. Summary.....	263
Are You Ready?.....	263
Further Reading.....	263
Help and Advice.....	264
How Can I Schedule My CCIE Lab Exam?.....	265
The Day Before.....	265
The Day of the Exam.....	265
Pass or Fail, What Next?.....	266

CCIE Routing and Switching v4.0 Configuration Practice Labs

Martin J. Duggan

ciscopress.com

About the Author

Martin James Duggan, CCIE No. 7942, is a network architect for AT&T. He designs network solutions for customers globally and specializes in data center networking and QoS. Martin mentors colleagues through their Cisco qualifications and holds regular internal training classes. Previous to this Martin was a network architect for IBM performing IP network designs and global network reviews. Martin has been in the industry for 20 years focusing on Cisco solutions for the previous 11 years. Martin is the co-author of the Cisco Press *CCIE Routing and Switching Practice Labs*, First Edition.

About the Technical Reviewer

Maurilio de Paula Gorito, CCIE No. 3807, is a triple CCIE, having certified in Routing and Switching in 1998, WAN Switching in 2001, and Security in 2003. Maurilio has more than 24 years of experience in networking, including Cisco networks and IBM/SNA environment. Maurilio's experience includes the planning, designing, implementation, and troubleshooting of large IP networks running RIP, IGRP, EIGRP, BGP, OSPF, QoS, and SNA worldwide. He also has more than 7 years of experience in teaching technical classes at schools and companies. Maurilio worked for Cisco as part of the CCIE team for 9 years. As the program manager for the CCIE Routing and Switching certification exams, Maurilio was responsible for managing the content development process for the CCIE Routing and Switching Lab and Written Exams, supporting candidates as part of the CCIE customer service, and proctoring CCIE lab exams at the CCIE lab in San Jose, CA, and worldwide. Maurilio also has presented Power Sessions at Cisco seminars and at CiscoLive. Maurilio currently works for Riverbed Technology as a certification manager responsible for overseeing the certifications and programs for Riverbed's Professional Services business unit. Maurilio is the co-author of the Cisco Press *CCIE Routing and Switching Practice Labs* and has reviewed several other Cisco Press books. Maurilio holds degrees in mathematics and pedagogy.

Dedication

Martin James Duggan: I would like to dedicate this publication to my family. Mum and Dad, thanks for your care and support in trying times recently for which I am extremely grateful. Neil and Jo, you are always there when I need your help. To my honorary CCNAs Anna and James, I am blessed to have children as wonderful as you. You are growing up far too quickly for my liking, but you make me the proudest father in the world.

Charlotte, what can I say? You are usually late but your timing when we met was impeccable; I cannot imagine you not being in my life now.

Acknowledgments

Martin James Duggan: This is my third opportunity to write for Cisco Press, so I would like to thank Brett Bartow for once again providing me with this enviable opportunity.

To Maurilio, who has reviewed this publication, I would like to say thank you for the time and experience you have put into this; you have shaped my work and I really value your contribution.

I'd like to thank my previous manager, Dave Mack. I was very lucky to have you as a manager Dave; you gave me some really interesting projects, encouraged me with this book, and were a pleasure to work with.

To Pete Davison and Mike (mountain goat) Jones, my cycling buddies who never seem to get bored with me talking networks or cracking Jethro jokes when we manage to get out, either that or they wanted me out of breath for the hills.

To Richard Burbage, my oldest friend, your suggestion really helped me, I owe you one.

Command Syntax Conventions

The conventions used to present command syntax in this book are the same conventions used in the IOS Command Reference. The Command Reference describes these conventions as follows:

- **Boldface** indicates commands and keywords that are entered literally as shown. In actual configuration examples and output (not general command syntax), boldface indicates commands that are manually input by the user (such as a **show** command).
- *Italics* indicate arguments for which you supply actual values.
- Vertical bars (|) separate alternative, mutually exclusive elements.
- Square brackets [] indicate optional elements.
- Braces { } indicate a required choice.
- Braces within brackets [{ }] indicate a required choice within an optional element.

Introduction

For more than ten years, the CCIE program has identified networking professionals with the highest level of expertise. Less than 3 percent of all Cisco certified professionals actually achieve CCIE status. The majority of candidates that take the exam fail at the first attempt because they are not fully prepared; they generally find that their study plan did not match what was expected of them in the exam. This practice exam has been designed to take you as close as possible to actually taking the real lab exam. It will show whether you are ready to schedule your lab, or if you need to reevaluate your study plan.

Exam Overview

The CCIE qualification consists of two exams, a 2-hour written exam followed by an 8-hour hands-on lab exam that now includes a troubleshooting section. Written exams are computer-based, multiple choice exams lasting 2 hours and available at hundreds of authorized testing centers worldwide. The written exam is designed to test your theoretical

knowledge to ensure you are ready to take the lab exam; as such, you are only eligible to schedule the lab exam after you have passed the written exam. Having purchased this publication, it is assumed that you have passed the written exam and are ready to practice for the lab exam. The lab exam is a 5 1/2-hour, hands-on exam in which you are required to configure a series of complex scenarios in strict accordance to the questions; it's tough but achievable. Troubleshooting is now included for 2 hours, and you are also presented with a series of further questions for a 30-minute period of the exam. Current lab blueprint content information can be found on the following URL:

<https://learningnetwork.cisco.com/docs/DOC-4603>.

Scoring Point System

In the actual exam a higher number of available points for certain questions would generally indicate that the required solution would take more time to achieve or that there would be multiple lines of configuration involved. This practice lab closely echoes the scoring system in place in the actual exam. If you find you are running short on time, try to get the smaller tasks completed and then return to the more complex questions.

Study Roadmap

Taking the lab exam is all about experience; you can't expect to take it and pass after just completing your written exam, relying on your theoretical knowledge. You will need to spend countless hours of rack time configuring features and learning how protocols interact with one another. To be confident enough to schedule your lab exam, review the following outlined points.

Assessing Your Strengths

Using the content blueprint, determine your experience and knowledge in the major topic areas. For areas of strength, practicing for speed should be your focus. For weak areas, you might need training or book study in addition to practice.

Study Materials

Choose lab materials that provide configuration examples and take a hands-on approach. Look for materials approved or provided by Cisco and its Learning Partners.

Hands-On Practice

Build and practice your lab scenarios on a per-topic basis. Go beyond the basics and practice additional features. Learn the **show** and **debug** commands along with each topic. If a protocol has multiple ways of configuring a feature, practice all of them.

Cisco Documentation CD

Make sure you can navigate the Cisco documentation CD with confidence because this is the only resource you will be allowed during the lab (or restricted access to the same content on Cisco.com). Make the CD part of your regular study; if you are familiar with it, you can save time during the exam.

Home Labs

Although acquiring a personal home lab is ideal, it can be costly to gather all the equipment you will need.

Cisco 360 Program

The Cisco 360 Learning Program encompasses six stages of activity to support successful learning for students:

1. **Assessment:** Students take a diagnostic pre-assessment lab to benchmark their knowledge of various networking topics.
2. **Planning:** Based on the pre-assessment, students create a learning plan that uses a mix of learning components to focus their study.
3. **Learning:** Students learn by participating in lessons and lectures, reading materials, and working with peers and instructors.

4. **Practice:** Students use the practice exercises to apply learning on actual network equipment.
5. **Mastery:** Students measure their understanding by completing assessments of knowledge and skill for various approaches to solving network problems.
6. **Review:** Students review their work with a mentor or instructor and tune their skills with tips and best practices.

Detailed information on the 360 program can be found on the following URL:

https://learningnetwork.cisco.com/community/learning_center/cisco_360/360-rs.

Equipment List and IOS Requirements

The lab exam tests any feature that can be configured on the equipment and the IOS versions indicated here:

- 1841 Series routers—IOS 12.4(T) – Advanced Enterprise Services
- 3825 Series routers—IOS 12.4(T) – Advanced Enterprise Services
- Catalyst 3560 Series switches running IOS version 12.2—Advanced IP Services

This page intentionally left blank

Practice Lab 1

The CCIE exam commences with 2 hours of troubleshooting followed by 5 1/2 hours of configuration and a final 30 minutes of additional questions. This lab has been timed to last for 8 hours of configuration and self-troubleshooting, so aim to complete the lab within this period. Then either score yourself at this point or continue until you feel you have met all the objectives. You will now be guided through the equipment requirements and prelab tasks in preparation for taking this practice lab.

If you don't own six routers and four switches, consider using the equipment available and additional lab exercises and training facilities available within the CCIE R&S 360 program. You can find detailed information on the 360 program and CCIE R&S exam on the following URLs, respectively:

https://learningnetwork.cisco.com/community/learning_center/cisco_360/360-rs

https://learningnetwork.cisco.com/community/certifications/ccie_routing_switching

Equipment List

You need the following hardware and software components to begin this practice lab:

- Six routers loaded with Cisco IOS Software Release 12.4 Advanced Enterprise image and the minimum interface configuration, as documented in Table 1-1

TABLE 1-1 Hardware Required per Router

Router	Model	Ethernet I/F	Serial I/F
R1	3825	1	1
R2	3725	1	2
R3	3825	1	1
R4	3825	2	—
R5	3825	2	1
R6	3825	2	—

NOTE

The 3825s used in this lab were loaded with `c3825-adventerprisek9-mz.124-6.T.bin`, and the 3725 was loaded with `c3725-adventerprisek9-mz.124-6.T.bin`.

NOTE

Notice in the initial configurations supplied that some interfaces will not have IP address pre-configured. This is because you either will not be using that interface or you need to configure this interface from default within the exercise. The initial configurations supplied should be used to preconfigure your routers and switch before the lab starts.

If your routers have different interface speeds than those used within this book, adjust the bandwidth statements on the relevant interfaces to keep all interface speeds in line. This can ensure that you do not get unwanted behavior due to differing IGP metrics.

- One 3550 switch with IOS 12.2 IP Services and three 3560 switches with IOS 12.2 IP Services

Setting Up the Lab

You can use any combination of routers as long as you fulfill the requirements within the topology diagram, as shown in Figure 1-1. However, it is recommended to use the same model of routers because this can make life easier if you load configurations directly from those supplied with your own devices.

Lab Topology

This practice Lab uses the topology outlined in Figure 1-1, which you need to re-create with your own equipment or by simply using the CCIE Assessor.

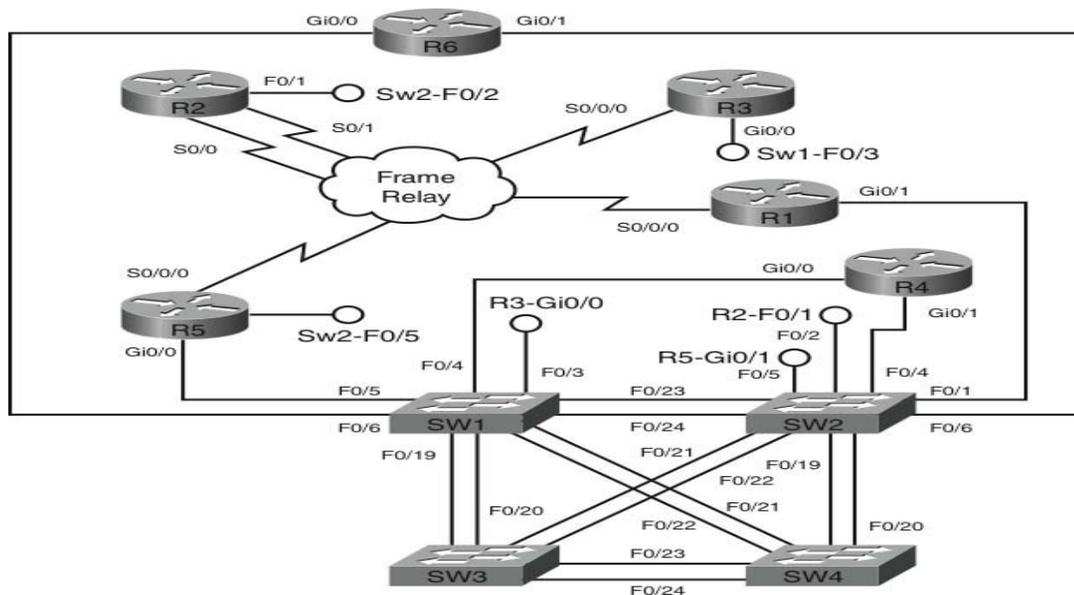


FIGURE 1-1
Lab Topology Diagram

Switch Instructions

Configure VLAN assignments from the configurations supplied or from Table 1-2 with the exception of Switch2 Fa0/4. (This will be configured during the lab.)

TABLE 1-2 VLAN Assignment

VLAN	Switch1	Switch2	Switch3	Switch4
34	Fa0/3, Fa0/4	—	—	—
45	Fa0/5	See Questions	—	—
46	Fa0/6	See Questions	—	—
100	—	Fa0/1	—	—
200	—	Fa0/2	—	—
300	I/F VLAN300	Fa0/5, Fa0/6, I/F VLAN300	I/F VLAN300	I/F VLAN300

Connect your switches with RJ45 Ethernet Cross Over cables, as shown in Figure 1-2.

NOTE

The CCIE Assessor topology version B is used for this lab. Additional interfaces available on the Assessor that are not required for this lab were omitted from Figure 1-1. If you are not using the CCIE Assessor, use Figure 1-1 and Figure 1-4 to determine how many interfaces you need to complete your own topology.

NOTE

Switch2 will be configured during the actual lab questions for VLAN45 and 46 interface Fa0/4.

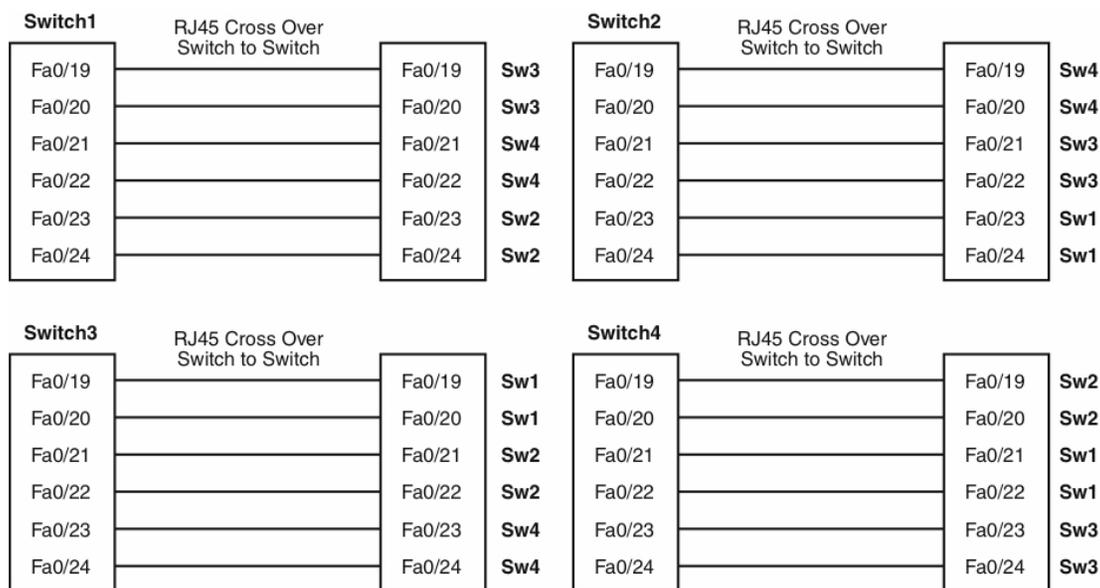


FIGURE 1-2 Switch to Switch Connectivity

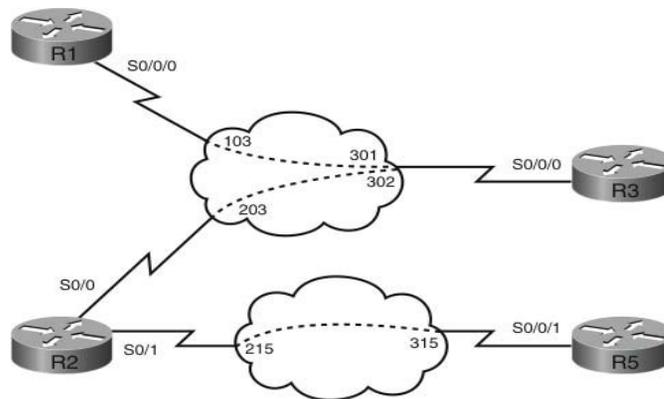
Frame Relay Instructions

Configure one of your routers you are going to use in the lab as a Frame Relay switch, or have a dedicated router purely for this task. This lab uses a dedicated router within the CCIE Assessor Version B topology for the Frame Relay switch. A fully meshed environment is configured between all the Frame Relay routers; pay attention in the lab as to which PVCs are actually required. Keep the encapsulation and Local Management Interface (LMI) settings to default for this exercise, but experiment with the settings outside the labs because you could be required to configure the Frame Relay switching within your actual lab.

If you are using your own equipment, keep the DCE cables at the frame switch end for simplicity and provide a clock rate to all links from this end.

The Frame Relay connectivity after configuration represents the logical Frame Relay network, as shown in Figure 1-3.

FIGURE 1-3
Frame Relay Logical
Connectivity



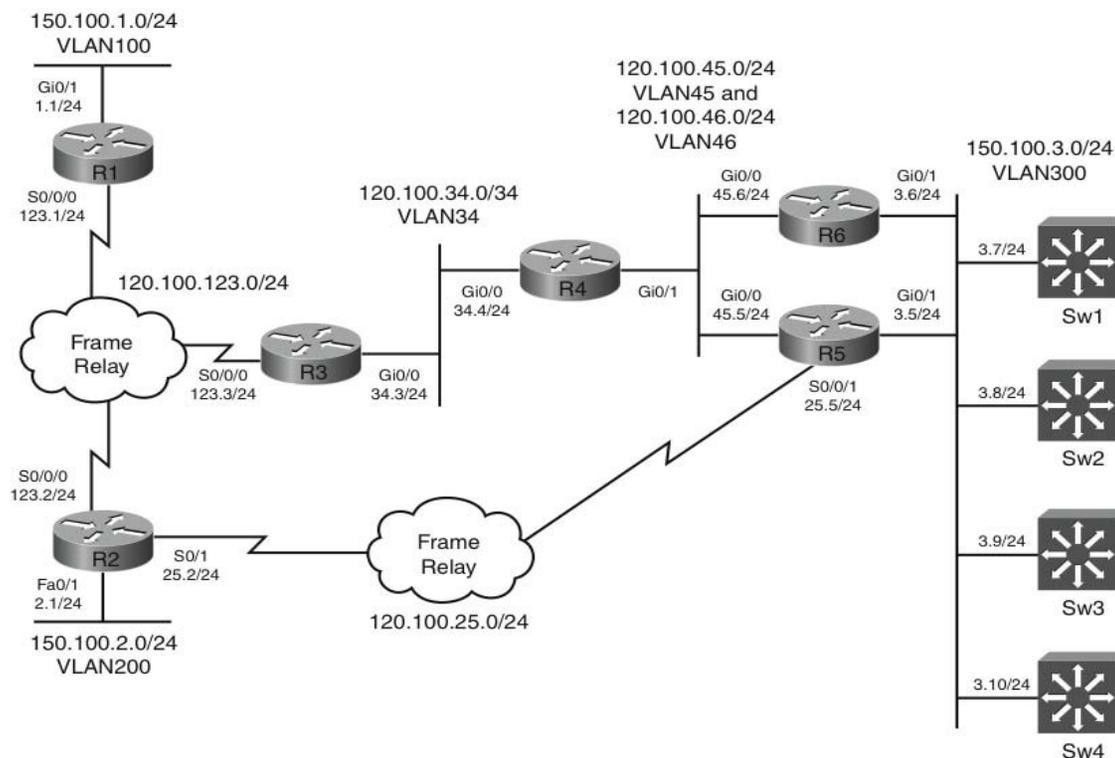
IP Address Instructions

You will find in the real CCIE lab that the majority of your IP addresses will be preconfigured; for this exercise you are required to configure your IP addresses, as shown in Figure 1-4, or load the initial router configurations supplied. If you are manually configuring your equipment, ensure you include the following Loopback addresses:

R1 Lo0 120.100.1.1/24	R6 Lo0 120.100.6.1/24
R2 Lo0 120.100.2.1/24	SW1 Lo0 120.100.7.1/24

R3 Lo0 120.100.3.1/24 SW2 Lo0 120.100.8.1/24
 R4 Lo0 120.100.4.1/24 SW3 Lo0 120.100.9.1/24
 R5 Lo0 120.100.5.1/24 SW4 Lo0 120.100.10.1/24

FIGURE 1-4
IP Addressing Diagram



Pre-lab Tasks

- Build the lab topology as per Figure 1-1 and Figure 1-2.
- Configure your Frame Relay switch router to provide the necessary Data Link Control Identifiers (DLCI) as per Figure 1-3.
- Configure the IP addresses on each router, as shown in Figure 1-4, and add the Loopback addresses. Alternatively, you can load the initial configuration files supplied if your router is compatible with those used to create

this exercise. R1 requires a secondary IP address on its GigabitEthernet 0/1 interface for this lab; details can be found on the accompanying initial configuration for R1.

General Guidelines

- Please read the whole lab before you start.
- Do not configure any static/default routes unless otherwise specified.
- Use only the DLCIs provided in the appropriate figures.
- Ensure full IP visibility between routers for ping testing/telnet access to your devices with exception to the Switch Loopback addresses. These will not be visible to the majority of your network because of the configuration tasks.
- If you find yourself running out of time, choose questions that you are confident you can answer; failing this choose questions with a higher point rating to maximize your potential score.
- Get into a comfortable and quiet environment where you can focus for the next 8 hours.
- Take a 30-minute break midway through the exercise.
- Have available a Cisco Documentation CD-ROM or access online the latest documentation from the following URL: http://www.cisco.com/en/US/products/ps6350/products_installation_and_configuration_guides_list.html.

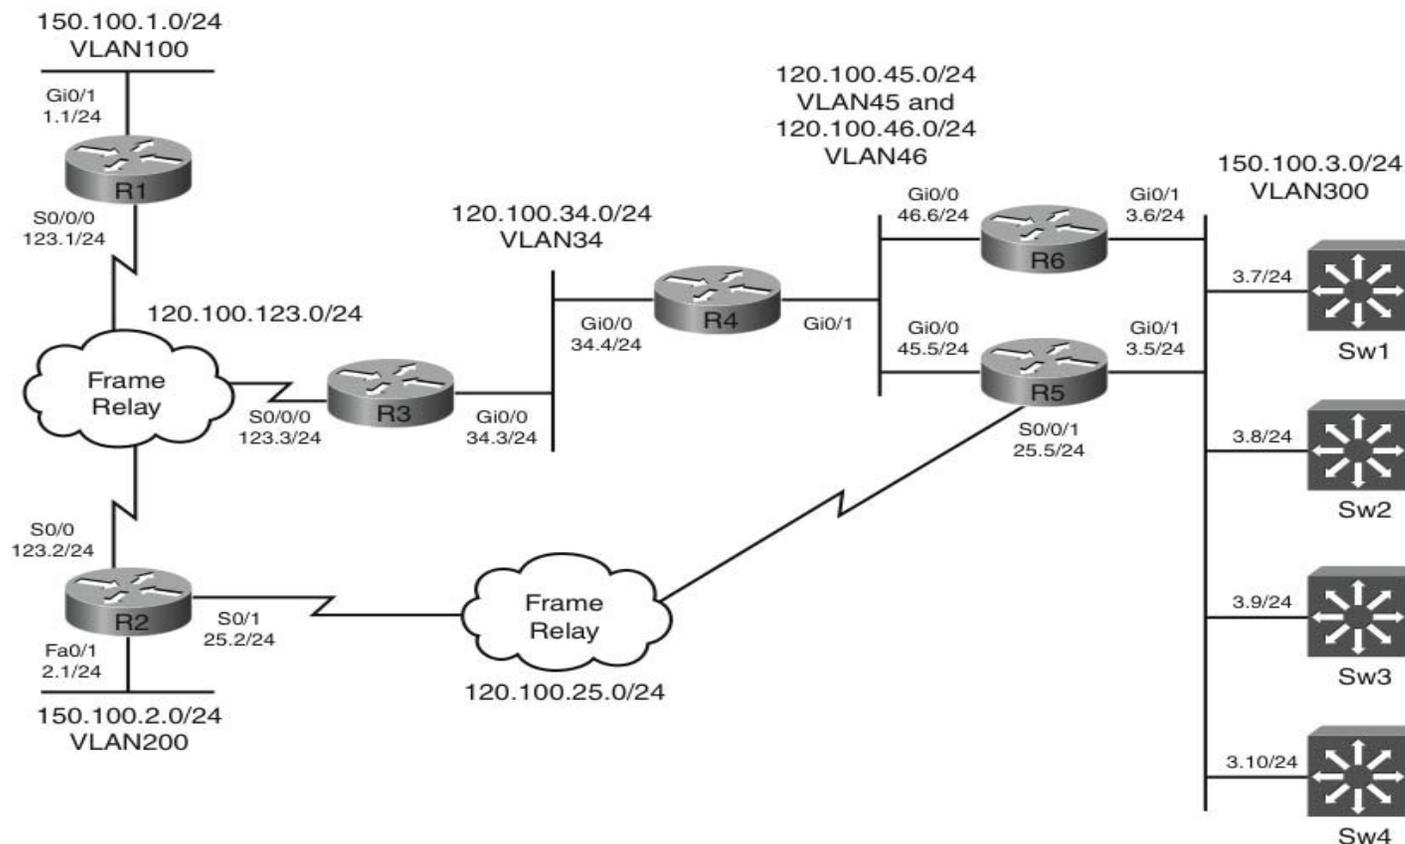
NOTE

Access only this URL, not the whole Cisco.com website; because if you are permitted to use documentation during your CCIE lab exam, it will be restricted. Consider opening several windows with the pages you are likely to look at to save time during your lab.

Practice Lab One

You will now answer questions in relation to the network topology, as shown in Figure 1-5.

FIGURE 1-5
Lab Topology Diagram

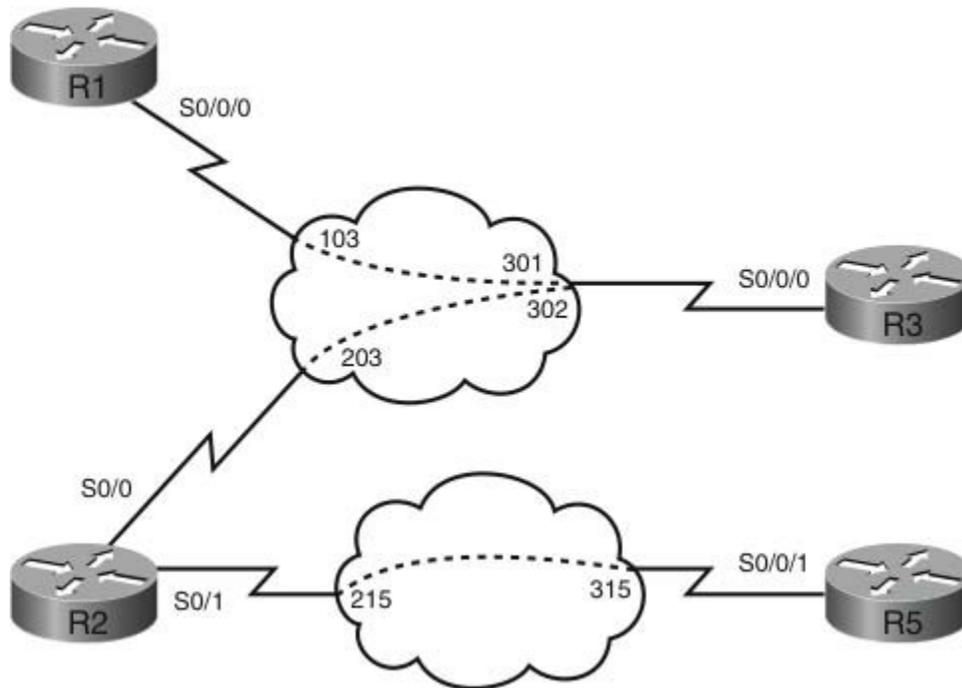


Section 1: LAN Switching and Frame Relay (28 Points)

- Configure your switches as a collapsed backbone network with Switches 1 and 2 performing core and distribution functionality and Switches 3 and 4 as access switches in your topology. Switches 3 and 4 should connect only to the core switches. (2 points)
- Switch 1 and 2 should run spanning tree in 802.1w mode; Switches 3 and 4 should operate in their default spanning-tree mode. (2 points)
- Configure Switch 1 to be the root bridge and Switch 2 the secondary root bridge for VLANs 1 and 300. Ensure that Switches 3 and 4 can never become root bridges for any VLANs for which Switch 1 and Switch 2 are root bridges by configuring only Switches 1 and 2. (2 points)

- Ensure you fully utilize the available bandwidth between switches by grouping together your interswitch links as trunks. Ensure that only dot1q and EtherChannel are supported. (3 points)
- Ensure traffic is distributed on individual Ethernet trunks between switches based on the destination MAC address of individual flows. (2 points)
- Ensure that user interfaces are shut down dynamically by all switches should they toggle excessively; if they remain stable for 35 seconds, they should be reenabled. Configure Fast Ethernet Port 0/10 on each switch so that if multicast traffic is received on this port, the port is automatically disabled. (2 points)
- Fast Ethernet Ports 0/11-17 will be used for future connectivity on each switch. Configure these ports as access ports for VLAN300, which should begin forwarding traffic immediately on connection. Devices connected to these ports will dynamically receive IP addresses from a DHCP server due to be connected to Port 0/18 on sw1. For security purposes, this is the only port on the network from which DHCP addresses should be allocated. Ensure the switches intercept the DHCP requests and add the ingress port and VLAN and switch MAC address prior to sending onward to the DHCP server. Limit DHCP requests to 600 packets per minute per user port. (6 points)
- For additional security ensure the user ports on Switches 1–4 and 11–17 can communicate only with the network with IP addresses gained from the DHCP feature configured previously. Use a dynamic feature to ensure the only information forwarded upon connection is DHCP request packets, then any traffic that matches the DHCP IP information received from the DHCP binding for additional security. (3 points)
- R5 and R6 have been preconfigured with IP addresses on their Ethernet interfaces. Configure R4 and its associated switch port accordingly without using secondary addressing to communicate with R5 and R6. Configure R4 with an IP address of 120.100.45.4/24 to communicate with R5, and configure R4 with an IP address of 120.100.46.4/24 to communicate with R6. Configure R4 Gi0/1 and Switch 2 FE0/4 only. (3 points)
- Your initial Frame Relay configuration has been supplied for the R1-R2-R3 connectivity and R2-R5. Configure each device per Figure 1-6 to ensure each device is reachable over the Frame Relay network. Use only the indicated DLCIs. (2 points)

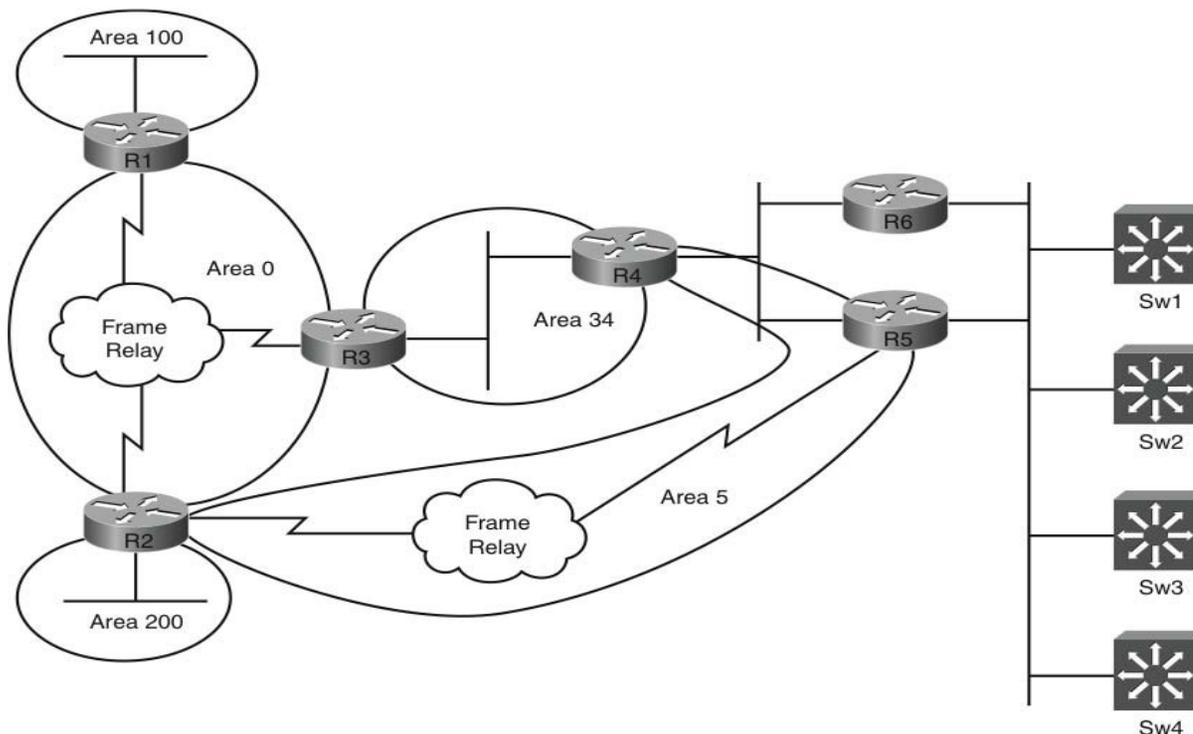
FIGURE 1-6
Frame Relay
Connectivity



Section 2: IPv4 IGP Protocols (22 Points)

Section 2.1: OSPF

FIGURE 1-7
OSPF Topology

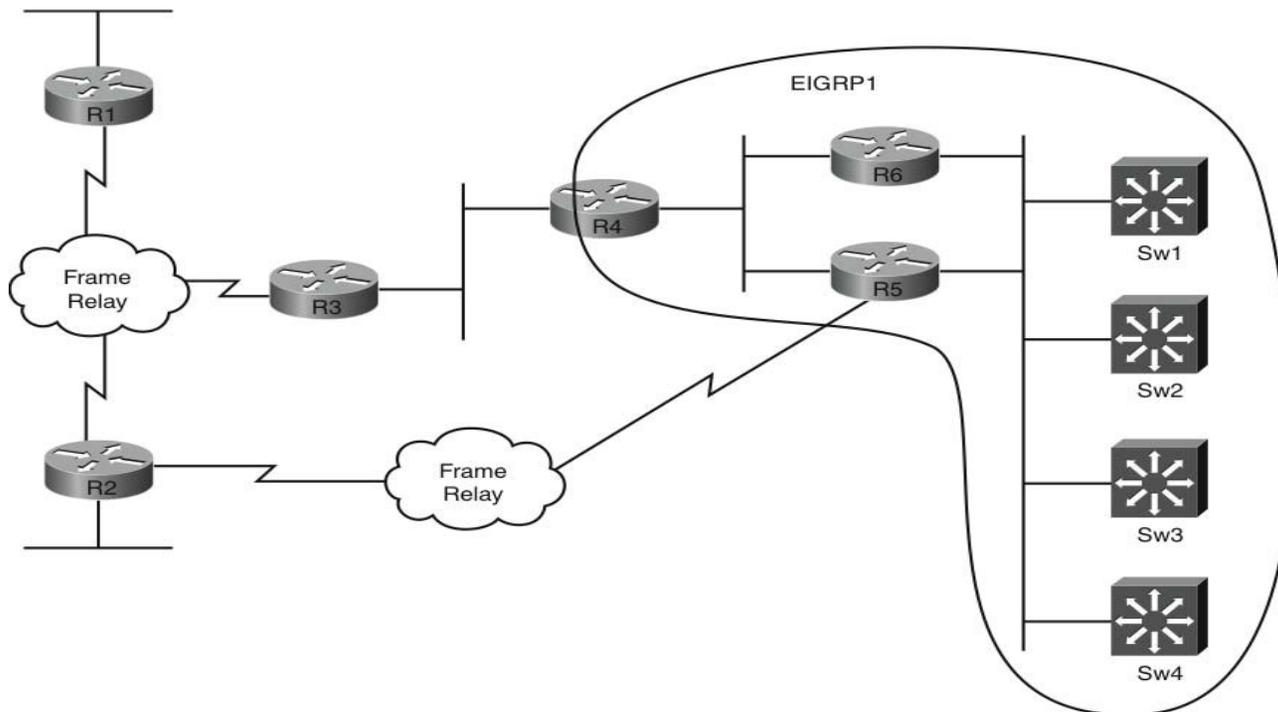


- Use a process ID of 1; all OSPF configuration where possible should not be configured under the process ID. Do not change the preconfigured interface types where applicable, The Loopback interfaces of Routers R1, R2, and R3 should be configured to be in Area 0. R4 should be in Area 34 and R5 in Area 5. (2 points)
- All Loopback networks should not be advertised as host routes. (1 point)
- Ensure that R1 does not advertise the preconfigured secondary address under interface Gigabit 0/1 of 120.100.100.1/24 to the OSPF network. Do not use any filtering techniques to achieve this. (2 points)

- R5 should use the Frame Relay link within Area 5 for its primary communication to the OSPF network. If this network should fail either at Layer 1 or Layer 2, R5 should form a neighbor relationship with R4 under Area 5 to maintain connectivity. Your solution should be dynamic ensuring that while the Area 5 Frame Relay link is operational there is no neighbor relationship between R4 and R5; however, the Ethernet interfaces of R4 and R5 must remain up. To confirm the operational status of the Frame Relay network, you should ensure that the serial interface of R5 is reachable by configuration of R5. You are permitted to define neighbor statements between R5 and R4. (4 points)

Section 2.2: EIGRP

FIGURE 1-8
EIGRP Topology



- Configure EIGRP using an AS number of 1. The Loopback interfaces of all routers and switches should be advertised within EIGRP. (2 points)

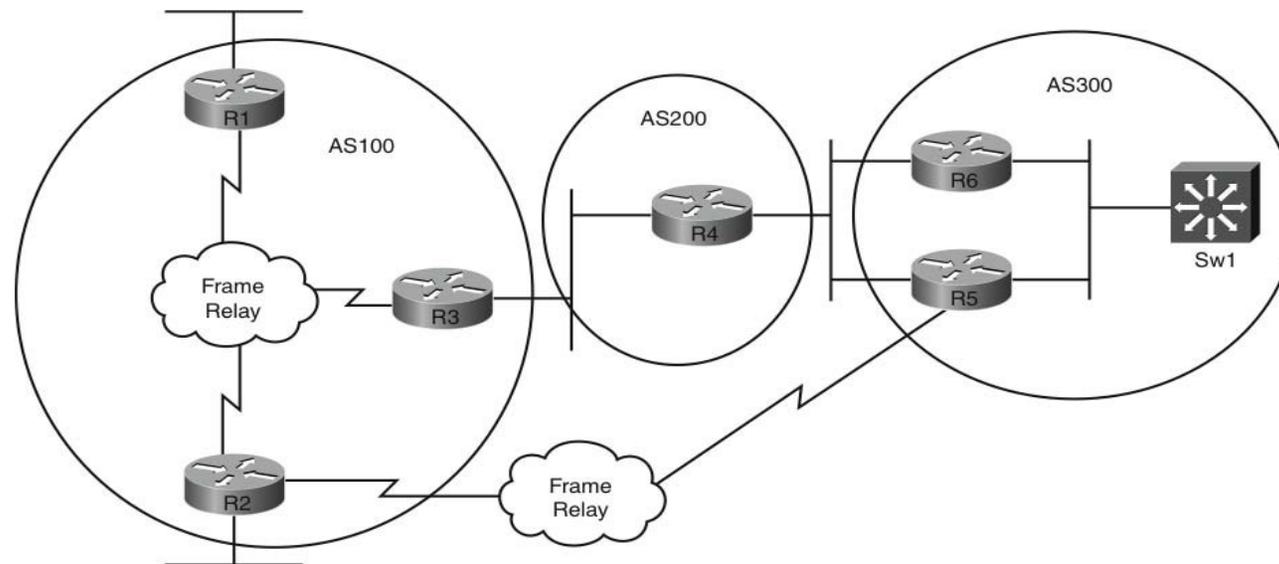
- Ensure that R4 does not install any of the EIGRP Loopback routes from any of the switches into its routing table; as such these routes should also not be present in the OSPF network post redistribution. Do not use any route-filtering ACLs, prefix lists, or admin distance manipulation to achieve this, and perform configuration only on R4. (3 points)
- R4 will have dual equal cost routes to VLAN300 (network 150.100.3.0) from R5 and R6. Ensure R4 sends traffic to this destination network to R5 rather than load sharing. If the route from R5 becomes unavailable, traffic should be sent to R6. You cannot policy route, alter the bandwidth, or delay statements on R4's interfaces, or use an offset list. Perform your configuration on R4 only. Your solution should be applied to all routes received from R5 and R6 as opposed to solely the route to network VLAN300. (3 points)

Section 2.3: Redistribution

- Perform mutual redistribution of IGP protocols on R4. All routes should be accessible with the exception of the switch Loopback networks because these should not be visible via R4 from an earlier question. EIGRP routes redistributed within the OSPF network should remain with a fixed cost of 5000 throughout the network. (3 points)
- Configure R4 to redistribute only up to five EIGRP routes and generate a system warning when the fourth route is redistributed. Do not use any access-lists in your solution. (2 points).

Section 3: BGP (14 Points)

FIGURE 1-9
BGP Topology

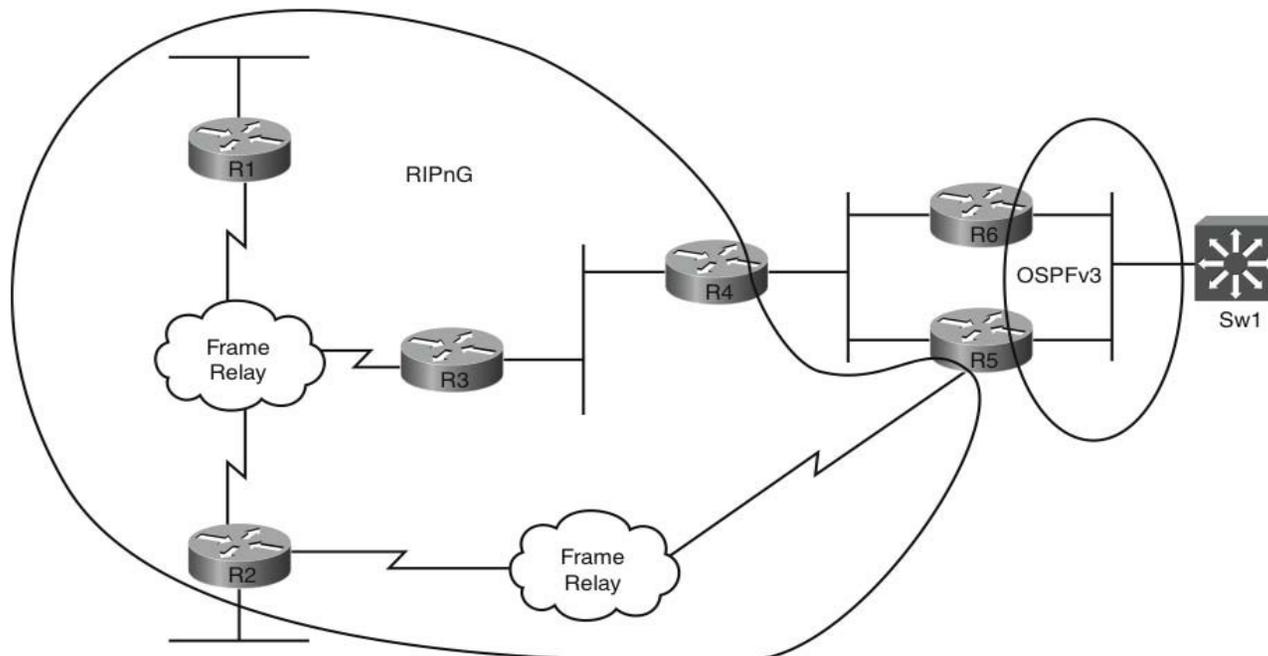


- Configure iBGP peering as follows: R1-R3, R2-R3, R6-R5, Sw1-R6, and Sw1-R5. Use minimal configuration and use Loopback interfaces for your peering. Configure eBGP peering as follows: R3-R4, R4-R6, R4-R5, and R5-R2. Use minimal configuration and use Loopback interfaces for your peering with the exception of R4 to R5. (2 points)
Use the AS numbers supplied in Figure 1-9. (2 points)
- AS200 is to be used as a backup transit network for traffic between AS100 and AS300; as such, if the FR network between R5 and R2 fails, ensure the peering between R2 and R5 is not maintained via the Ethernet network. Do not use any ACL type restrictions or change the existing peering. (2 points)
- Configure a new Loopback interface 2 on R2 of 130.100.200.1/24, and advertise this into BGP using the network command. Configure R2 in such a way that if the Frame Relay connection between R2 and R5 fails, AS300 no longer receives this route. Do not use any filtering between neighbors to achieve this or neighbor-specific commands. (3 points)

- Configure HSRP between R5 and R6 on VLAN300 with R5 active for .1/24. If the network 130.100.200.0/24 is no longer visible to AS300, R6 should dynamically become the HSRP active. Configure R5 to achieve this solution. (4 points)
- Configure two new Loopback interfaces on R1 and R2 of 126.1.1.1/24 and 130.1.1.1/24, respectively, and advertise these into BGP using the **network** command. R3 should be configured to enable only BGP routes originated from R1 up to network 128.0.0.0 and from above network 128.0.0.0 originated from R2. Use only a single ACL on R3 as part of your solution. (3 points)

Section 4: IPv6 (14 Points)

FIGURE 1-10
IPv6 Topology



- Configure IPv6 addresses on your network as follows:
2007:C15:C0:10::/64 – R1 Gi0/0
2007:C15:C0:11::/64 – R1 S0/0/0

2007:C15:C0:11::2/64 – R2 S0/0
2007:C15:C0:11::3/64 – R3 S0/0/0
2007:C15:C0:12/64 - R2 FE0/1
2007:C15:C0:14::2/64 – R2 S0/1
2007:C15:C0:14::5/64 – R5 S0/0/1
2007:C15:C0:15::3/64 – R3 Gi0/0
2007:C15:C0:15::4/64 – R4 Gi0/0
2007:C15:C0:16::5/64 – R5 Gi0/1
2007:C15:C0:16::6/64 – R6 Gi0/1

Section 4.1: RIPng

- Configure RIPng ensuring your IPv6 routes are visible throughout your RIPng domain. Do not disable split-horizon. (3 points)

Section 4.2: OSPFv3

- Configure OSPFv3 with a process ID of 1 with all OSPF interfaces assigned to Area 0. (2 points).
- The IPv6 network is deemed to be stable; therefore, reduce the number of LSAs flooded within the OSPF domain. (2 points)

Section 4.3: Redistribution

- Redistribute RIPng routes into the OSPFv3 domain (one way). RIP routes should have a fixed cost of 5000 associated to them within the OSPF network. (1 point)
- Ensure the OSPFv3 network is reachable from the RIP network by a single route of 2007::16, which should be seen within the RIP domain. Configure R5 only to achieve this. The OSPF domain should continue to receive specific RIPng subnets. (2 points)
- Ensure that if the serial link fails between the OSPF and RIPng domain, routing is still possible between R5 and R4 over VLAN45. Do not enable RIP on the VLAN45 interfaces of R4 and R5. Configure R4 and R5 to achieve this, which should be considered as an alternative path only if a failure occurs. (3 points)

- Ensure the summary route configured previously is not seen back on the routing table of R5; configure only R5 to achieve this. (1 point)

Section 5: QoS (8 Points)

- You are required to configure QoS on switch1 according to the Cisco QoS baseline model. Create a Modular QoS configuration for all user ports (Fast Ethernet 1-24) that facilitates the following requirements (3 points):
 - 1) All ports should trust the DSCP values received from their connecting devices.
 - 2) Packets received from the user ports with DSCP values of 48, 46, 34, 32, 24, 28, 16, and 10 should be remarked to DSCP 8 (PHB CS1) in the event of traffic flowing above 5 Mbps on a per port basis. This traffic could be a combination of any of the preceding DSCP values with any source/destination combination. Ensure a minimum burst value is configured above the 5 Mbps.
- Switch1 will be connected to a new trusted domain in the future using interface gigabit 0/1. A DSCP value received locally on sw1 of AF43 should be mapped to AF42 when destined for the new domain. (2 points)
- Configure Cisco Modular QoS as follows on R1 for the following traffic types based on their associated Per Hop Behavior into classes. Incorporate these into an overall policy that should be applied to the T1 interface S0/0/0. Assume a PVC of line rate on the Frame Relay network and allow each class the effective bandwidth as detailed (2 points):

Class	PHB	Assigned Speed
Routing	CS6	46 Kbps
VoIP	EF	247 Kbps
Interactive Video	AF41	247 Kbps
Mission Critical Data	AF31	247 Kbps
Call-Signaling	CS3	46 Kbps
Transactional Data	AF21	216 Kbps
Network-Mgmt	CS2	46 Kbps
Bulk Data	Af11	46 Kbps
Scavenger	CS1	15 Kbps
Default	0	386 Kbps

- Configure R2 so that traffic can be monitored on the Frame Relay network with a view to a dynamic policy being generated in the future that trusts the DSCP value of traffic identified on this media. (1 point)

Section 6: Security (6 Points)

- Configure R3 to identify and discard the following custom virus; the virus is characterized by the ASCII characters “Hastings_Beer” within the payload and utilizes UDP ports 11664 to 11666. The ID of the virus begins on the third character of the payload. The virus originated on VLAN 34. (2 points)
- An infected host is on VLAN 200 of 150.100.2.100; ensure that only within BGP AS10, traffic destined for this host is directed to null0 of each local router. You cannot use any ACLs to block traffic to this host specifically but can use a static route pointing to null 0 for traffic destined to 192.0.2.0 /24 on routers within AS10. R2 can have an additional static route pointing to null0. Use a BGP feature on R2 to ensure traffic to this source is blocked. Prevent unnecessary replies when traffic is passed to the null0 interface for users residing on VLAN100. (3 points)
- In a view of protecting the control plane on Router R6, configure CoPP so that IP Packets with a TTL of 0 or 1 are dropped rather than processed with a resulting ICMP redirect sent to the originator. (1 point)

Section 7: Multicast (4 Points)

- Configure Routers R1, R2, R3, and R4 for IPv4 Multicast; configure R3 to send multicast advertisements of its own time by use of NTP sourced from interface Gig 0/0. Configure PIM sparse mode on all required interfaces. R3 should also be used to advertise its own gigabit interface IP address as an RP. R3 should also advertise the IP address you are using for the NTP advertisements that will be 224.0.1.1. Do not use the command **ntp server** in any configurations. Routers R1, R2, and R4 should all show a clock synchronized to that of R3. (4 points)

IP Services (4 Points)

- Configure the following commands on Router R1:
aaa new-model
logging buffered
logging 120.100.99.1

Configure a policy on Router R1 so that if a user tries to remove AAA services or disable logging via the CLI that a syslog message of UNAUTHORIZED-COMMAND-ENTERED is generated. The policy should ensure either command is not executed and should consist of a single-line command for the CLI pattern detection. The policy and CLI should run asynchronously. The policy should also generate an email from the router to a mail server residing on IP address 120.100.99.2 (to security@lab-exam.net from eem@lab-exam.net subject "User- Issue" with the message body consisting of details of who was logged on the time either of the commands were entered). (2 points)

Cisco WAAS devices are to be installed on Switches 1 and 2 in the future on VLAN300. Configure Routers R5 and R6 to provide WCCPv2 redirection for clients residing on VLAN300 to ensure that all TCP traffic other than telnet is redirected only to the WAEs that will reside on addresses 150.100.3.50 and .51 within VLAN300. You are not required to configure the switches for WCCP and can assume that incoming WAAS traffic from the network will arrive at interfaces Gi0/0 on both R5 and R6. Secure your WCCP with this password: CCIE. (2 points)

“Ask the Proctor”

NOTE

This section should be used only if you require clues to complete the questions. In the actual CCIE lab, the Proctor will not enter into any discussions regarding the questions or answers; he or she will be present to ensure you do not have problems with the lab environment and to maintain the timing element of the exam.

Section 1: LAN Switching and Frame Relay

Q: Do you want me to configure the collapsed backbone network by manipulating spanning tree to ensure that Switch 1 and Switch 2 are the cores for each VLAN in use?

A: You are requested to configure root bridges in a later question.

Q: All the switches are already connected, so I can't change this unless I shut down some of the connections between switches. Is this acceptable?

A: Yes.

Q: If I explicitly configure Switches 1 and 2 as root bridges, surely this will never enable Switches 3 and 4 to become root bridges?

A: No it won't. If a superior BPDU is received on ports connecting to Switches 3 and 4 from Switches 1 and 2, Switches 3 and 4 could become root bridges; use a feature that effectively ignores a superior BPDU if received.

Q: Do you want me to disable spanning tree down to Switches 3 and 4? Is this acceptable?

A: No, spanning tree must remain in operation.

Q: Can I configure a MAC address type access-list to block all multicast at Layer 2?

A: No, this wouldn't disable the port if multicast traffic was present on it; look for a dynamic solution that does not require an ACL.

Q: Can I configure the **switchport block multicast** command?

A: No, this would block the traffic but wouldn't disable the port.

Q: Would you like me to VLAN load balance to utilize bandwidth?

A: No, the question directs you how to use the trunks.

Q: Would you like me configure Switch 1 to allocate DHCP addresses?

A: No, the question relates to a fictitious DHCP server that would be connected to Fa0/18 on Switch1.

Q: Can I manipulate a helper-address function to answer the DHCP question by using ACLs?

A: No, use a recognized DHCP security-related solution.

Q: Can I configure port security to bind my MAC addresses?

A: No, use a feature that complements your DHCP solution.

Q: Can I just configure R4 to trunk to Switch2 and have a subinterface in both VLAN45 and VLAN46?

A: Yes.

Q: I've configured my trunk on Switch2 to R4 and I can't ping between R4 and R5; similarly I can't ping between R4 and R6. Is there anything else I need to do?

A: Remember the switches are in VTP transparent mode; you might want to check that Switch2 has the required VLANs configured to enable propagation within your switched network.

Q: My Frame Relay network picks up the DLCIs automatically. is this okay?

A: No, you need to ensure that you do not use additional DLCIs other than those specified.

Q: Do you want me to manually map to the DLCIs I should be using?

A: Yes.

Section 2: IPv4 IGP Protocols

Section 2.1: OSPF

Q: I am used to configuring OSPF under the process; surely this is the only place I can configure the parameters?

A: There have been recent advances in OSPF enabling you to configure it purely under specific areas of the router rather like with IPv6. Take a look at the commands available to you under the interfaces.

Q: My neighbor relationship is down over the Frame Relay network. I notice I have different OSPF network types pre-configured. Can I change these?

A: No, use an alternative method of bringing the interface parameters back into line.

Q: My secondary address is advertised automatically under OSPF; can I use a distribute-list or prefix type list to block it?

A: No, use an OSPF feature to disable the advertisement of this secondary address.

Q: I've attempted to form a neighbor relationship with R4 from R5 using a backup interface. Is this okay?

A: No, the question states that your solution should cater for either Layer 1 or Layer 2 failures and that the Ethernet should remain up. Backup interfaces would be fine for a Layer 1 failure but not for a Layer 2 type issue if you had problems with specific DLCIs that caused neighbor failures over the Frame relay. This feature would also ensure the Ethernet network would be down until the backup interface is activated.

Q: How about an OSPF demand circuit between R4 and R5?

A: No, this would involve a neighbor relationship being maintained. You need to allow only the neighbor relationship to be formed if a failure condition occurs.

Q: Can I use BFD between R4 and R5?

A: No, this might aid in failure detection, but it does not meet the objectives of the question.

Q: To confirm the operation status of R5's serial interface, can I just ping it?

A: You can use ICMP but you need to ensure your solution is dynamic.

Q: My Frame Relay is up on R5 and I can ping across it to R2 from R5, but I can't ping my own Frame Relay interface. Is this normal?

A: Yes, perform a debug of the Frame Relay packets if you need to; remember what you need to gain IP connectivity on a Frame Relay network.

Q: If I use IP SLA to automatically ping R5 to check the status, is this okay?

A: Yes.

Q: Okay, I have IP SLA running but I'm stuck. Is this anything to do with tracking the response to the ping?

A: Yes.

Q: How about if I use policy routing with the next hop based on the tracking status?

A: This is fine; just remember that this traffic will be based locally on the router when applying any policies.

Q: I've worked out how to do this and managed to get a neighbor up when the Frame Relay fails, but my OSPF connectivity is still not perfect through the Ethernet. Is this normal?

A: Not if you have configured correctly; take a look at your topology and areas. Something might have changed when R5 connects over the Ethernet.

Section 2.2: EIGRP

Q: If I advertise my Loopbacks into EIGRP won't that mean that R4 and R5 will have their Loopbacks advertised by both OSPF and EIGRP?

A: Yes, this is fine and is in accordance with the question.

Q: To stop R4 from receiving the Switch Loopbacks can I stop advertising them from the switches?

A: No, you should use a feature on R4 to block them.

Q: Can I use a neighbor prefix list to block the Loopbacks?

A: No, you cannot use any type of ACLs or prefix lists.

Q: I've noticed when I look at the specific Loopback routes that they have a hop count associated with them. It's unusual to associate hop counts with EIGRP, but can I block routes based on their hop count?

A: Yes.

Q: If I can't change the bandwidth and delay on R4, can I use a route-map to manipulate the EIGRP K values associated on a per neighbor basis?

A: Yes.

Section 2.3: Redistribution

Q: Do you require a distribute-list to block the switch Loopbacks from entering the OSPF domain?

A: No, you should have blocked these from entering your IP routing table within R4 previously, so additional blocking would not be required.

Q: I have only one redistribution point, and there is no benefit in creating filtering to protect against potential routing loops between protocols. Is this acceptable?

A: Yes, in this scenario this would be superfluous.

Q: Can I use a route-map to enable five specific EIGRP routes to be redistributed into OSPF?

A: No, the question doesn't guide you to redistribute specific routes. Use a more general method of allowing a specific number of routes.

Sect

Q: Is it okay to disable auto synchronization in BGP?

A: You need to determine whether you need this feature on or off. Remember that you should have synchronization on only when you are fully redistributing between BGP and your IGP.

Q: Do you want me to configure ebgp multihop but limit it to a value of 2 on R3 for a TTL security check?

A: There is a specific security configuration feature within BGP to perform the TTL check.

Q: If I use the TTL security hops with a value of 2, is this all you are looking for?

A: You need to ensure that your peering still works effectively between R3 and R4 when you have configured this feature.

Q: I find that when the Frame Relay network fails my neighbor relationship is still maintained between R2 and R5. This is because the Loopback routes are still available over the alternative path through the network. Can I block my Loopbacks or policy route at some point to effectively break the peering?

A: You do need to effectively break the peering, but there is a far simpler method of achieving this that still maintains unaltered communication between R2 and R5. Think about what you need to configure when you have EBGP peers.

Q: I might have been a little generous with my original multihop value between R2 and R5. If I reduce this to a TTL of 2, I can break the peering. Is this okay?

A: Yes.

Q: I think I can stop the Loopback on R2 being advertised by using the community value of no-export, but if I enable this to R2, it wouldn't make to R5 even when the Frame Relay is working?

A: Correct, it wouldn't be advertised to R5 AS300 from R2. Just think about whether R2 is the best place to send the community to originally.

Q: For the HSRP question is this some form of conditional advertising?

A: No, the clue is in the question; just find a way of tracking the BGP route and manipulate the HSRP process.

Q: If I enable IP SLA to track a route in the routing table, can I use this to control HSRP?

A: Yes.

Q: You haven't told me what address I should use for HSRP. Is it okay to use the first address in the subnet?

A: Yes.

Q: I have configured my two new Loopbacks; can I use two route-maps inbound from R1 and R2 both pointing to different ACLs so that each route-map calls only one ACL?

A: No, you still have two ACLs.

Q: Can I set community values on the routes and match on these using a single ACL?

A: No, you are instructed to use an ACL; your solution would require additional configuration.

Q: Can I use a prefix-list to achieve this?

A: No, you are instructed to use an ACL.

Q: So I need an ACL with a mask suitable for both ranges?

A: Not necessarily; you would need to match only one requirement on the permit functionality; the other could be met by deny.

Section 4: IPv6

Q: Should I use the eui-64 address format when configuring my addresses?

A: No, if these were required you would have been instructed to do so in the question.

Q: I've configured my IPv6 addresses and created a Frame Relay map for these on my existing DLCIs but still can't ping across the Frame Relay network. Should I be able to?

A: Yes, if you debug your Frame Relay traffic, you will find you need additional configuration.

Q: I have configured RIPng between R1, R3, and R2; R3 receives both spoke routes but R1 does not see the R2 IPv6 route and vice versa. If this is split-horizon behavior and I can't disable it, can I create subinterfaces on my Frame Relay network?

A: No, use a feature that is common when running IPv6 over IPv4 networks.

Q: Can I tunnel between R1 and R2?

A: Yes.

Q: You are not requesting mutual redistribution between RIPng and OSPFv3. How will my RIPng domain communicate with the OSPFv3 domain?

A: This issue is addressed in the following task.

Q: If I can't use RIPng directly on VLAN45 between R4 and R5, can I configure OSPFv3 on VLAN45?

A: No, find a way to still run RIPng between routers without enabling it on the physical interfaces.

Q: Can I tunnel between R4 and R5?

A: Yes.

Section 4.3: Redistribution

Q: I have redistributed RIPng into OSPFv3 on R5, which is the only suitable location, and noticed that in my OSPFv3 domain I do not see the IPv6 network configured on the Frame Relay network between R2 and R5. Is this okay?

A: No, this network should be advertised to the OSPFv3 domain. Use a feature within the OSPFv3 process as you would to overcome this if this were IPv4 redistribution.

Q: Can I redistribute a static IPv6 route on R5 into RIPng for 2007::/16?

A: No, static routes are permitted unless specified. What would you do if this were IPv4?

Q: If I can't enable RIPng on VLAN45 between R4 and R5, can I enable OSPFv3?

A: No, this would also require you to perform redistribution at this point?

Q: How about tunneling again and enabling RIPng over the tunnel. Is this OK ?

A: Yes.

Q: I have created my tunnel and found that this is now the primary route rather than an alternative path. Can I perform some kind of backup interface to make this come up only if a failure occurs on the Frame Relay?

A: No, you haven't been given sufficient information to make this judgment. This approach would also break your IPv4 network; think why the Ethernet path is preferred and manipulate it.

Q: Can I use a prefix-list to block the summary and permit all other IPv6 routes?

A: Yes, this is fine.

Section 5: QoS

Q: Can I just trust DSCP on my physical ports?

A: No, this should be completed as part of your policy.

Q: Shall I rate-limit my ports to 5M on a per-port basis?

A: No, this should be completed as part of your policy.

Q: You haven't indicated what the minimum burst size should be, is this correct?

A: Yes, just use the available limits within the command options.

Q: I believe I can use a DSCP mutation map to convert the DSCP values for the future, but the command won't take the values AF43 and AF42.

A: No, it won't because these are Assured Forwarding values. You need to convert these to DSCP values; search your Documentation CD or available Cisco.com pages.

Q: I am trying to assign bandwidth within my class with the speeds supplied, but I can see only a percentage option, is this correct?

A: Yes, you need to do some math. You are supplied with the information you require and just need to remember how fast a T1 line is.

Section 6: Security

Q: Can I use a route-map and ACLs to identify the traffic by port number?

A: No, this would identify the UDLD traffic but not the virus payload as per the question. Investigate the options open to you with NBAR.

Q: Can I policy route traffic destined to the infected host to null0?

A: No, you need to use a BGP-related feature.

Q: A static route for 192.0.2.0/24 won't have any bearing on traffic destined to the infected host, why is this relevant?

A: Think about the way BGP works. It's the only routing protocol where you don't need to be directly connected to form a neighbor relationship; as such you transport next-hop information with your updates.

Q: I have configured CoPP on R6 and seem to have lost all my routes. Is this expected behavior? Do you want me to fix this as part of the CoPP question?

A: If you have lost your routes, think about why this has happened. Yes, provide a fix otherwise you would lose points in other sections.

Section 7: Multicast

Q: If I can't configure `ntp server` on R1, R2, and R4, there won't be a way I can get these routers to peer with R3. Is this correct?

A: Yes, you don't need to specifically peer with R3 as the server; you should aim to receive the ntp stream though that R3 should be configured to multicast.

Q: Do you want me to create and announce the group 224.0.1.1 on R3?

A: Yes.

Section 8: IP Services

Q: I guess this is an EEM question looking at the email address?

A: Correct.

Q: Do you need me to set up a route to 120.100.99.0/24?

A: No.

Q: I can't get both commands onto a single CLI pattern event. Is it okay to configure two?

A: No, you are directed to configure a single CLI pattern event command that will pick up either command.

Q: Do you want a GRE type redirection for the WCCP?

A: No, you have not been given sufficient information for GRE mode, or indeed if you should configure tunnels and so on; keep your configuration simple and follow the question.

Q: Should I block telnet and then permit all other IP traffic?

A: Think about what WAAS achieves. does it optimize all IP traffic or just specific protocols?

Q: Should I configure WCCP services 61 and 62 on the switches for VLAN300?

A: No, you are directed to configure only the routers.

Lab Debrief

The lab debrief section now analyzes each question showing you what was required and how to achieve the desired results. You should use this section to produce an overall score for this practice lab.

Section 1: LAN Switching and Frame Relay (28 Points)

- Configure your switches as a collapsed backbone network with Switches 1 and 2 performing core and distribution functionality and Switches 3 and 4 as access switches in your topology. Switches 3 and 4 should connect to only the core switches. (2 points)

This is a simple start to the exercise. The switches are fully meshed to begin with; to create a collapsed backbone topology, the core switches should be connected together, and each access switch should be dual-homed to the core switches. The only switches that should not connect directly to each other would be the access switches (Sw3 and Sw4). By shutting down the interfaces between Sw3 and Sw4, you create the required topology. If you have configured this correctly, as shown in Example 1-1, you have scored 2 points. Even though the resulting topology is not looped at this stage, you can verify route bridge assignment by using the **show spanning tree root** command.

EXAMPLE 1-1 Sw3 and Sw4 Configuration

```
SW3(config)# interface range fastEthernet 0/23-24
```

```
SW3(config-if-range)# shut
```

```
SW4(config)# interface range fastEthernet 0/23-24
```

```
SW4(config-if-range)# shut
```

- Switch 1 and 2 should run spanning tree in 802.1w mode. Switches 3 and 4 should operate in their default spanning-tree mode. (2 points)

802.1w is rapid spanning tree; this is backward compatible with the switches' default (PVST), so by configuring Switches 1 and 2 into rapid spanning tree mode, spanning tree can still operate effectively with Switches 3 and 4. If you have configured this correctly, as shown in Example 1-2, you have earned another 2 points.

EXAMPLE 1-2 Sw1 and Sw2 Configuration

```
SW1(config)# spanning-tree mode rapid-pvst
```

```
SW2(config)# spanning-tree mode rapid-pvst
```

- Configure Switch 1 to be the root bridge and Switch 2 the secondary root bridge for VLANs 1 and 300. Ensure that Switches 3 and 4 can never become root bridges for any VLANs for which Switch 1 and Switch 2 are root bridges by configuring only Switches 1 and 2. (2 points)

This is a straightforward question for the core switches. The root bridge prioritization root guard is configured on the ports that connect Switches 1 and 2 to Switches 3 and 4; this ensures that if a superior BPDU is received on these ports, it is ignored. If you have configured this correctly, as shown in Example 1-3, you have 2 points.

EXAMPLE 1-3 Sw1 and Sw2 Root Bridge Configuration

```
SW1(config)# spanning-tree vlan 1 root primary
SW1(config)# spanning-tree vlan 300 root primary
SW1(config-if)# interface FastEthernet 0/19
SW1(config-if)# spanning-tree guard root
SW1(config-if)# interface FastEthernet 0/20
SW1(config-if)# spanning-tree guard root
SW1(config-if)# interface FastEthernet 0/21
SW1(config-if)# spanning-tree guard root
SW1(config-if)# interface FastEthernet 0/22
SW1(config-if)# spanning-tree guard root
```

```
SW2(config)# spanning-tree vlan 1 root secondary
SW2(config)# spanning-tree vlan 300 root secondary
SW2(config-if)# interface FastEthernet 0/19
SW2(config-if)# spanning-tree guard root
SW2(config-if)# interface FastEthernet 0/20
SW2(config-if)# spanning-tree guard root
SW2(config-if)# interface FastEthernet 0/21
SW2(config-if)# spanning-tree guard root
SW2(config-if)# interface FastEthernet 0/22
SW2(config-if)# spanning-tree guard root
```

- Ensure that you fully utilize the available bandwidth between switches by grouping your interswitch links as trunks. Ensure that only dot1q and EtherChannel are supported. (3 points)

This is another straightforward question for all switches to create EtherChannels between devices. Using the command **channel-group *n* mode on** under the physical interfaces ensures that only EtherChannel is supported, as opposed to pagp or lacp, and dot1q is the trunking protocol. For Layer 2 EtherChannels, you don't have to create a port-channel interface first by using the **interface port-channel** configuration command before assigning a physical port to a channel group. You can use the **channel-group** interface configuration command that automatically creates the port-channel interface, although a manual port channel configuration has been shown here for clarity. Remember that now that you have EtherChannels between switches, you will need to configure root guard on these interfaces to ensure that Switches 3 and 4 cannot become root bridges. This is over and above the previous physical interface configuration completed previously. If you have configured this correctly, as shown in Example 1-4, you have scored 3 points.

EXAMPLE 1-4 Switch 1, 2, 3, and 4 EtherChannel Configuration

```
SW1(config)# interface Port-channel1
SW1(config-if)# switchport trunk encapsulation dot1q
SW1(config-if)# switchport mode trunk
SW1(config-if)# spanning-tree guard root
SW1(config-if)# interface Port-channel2
SW1(config-if)# switchport trunk encapsulation dot1q
SW1(config-if)# switchport mode trunk
SW1(config-if)# spanning-tree guard root
SW1(config-if)# interface Port-channel3
SW1(config-if)# switchport trunk encapsulation dot1q
SW1(config-if)# switchport mode trunk
SW1(config-if)# interface range FastEthernet0/19-20
SW1(config-if)# channel-group 1 mode on
SW1(config-if)# interface range FastEthernet0/21-22
SW1(config-if)# channel-group 2 mode on
SW1(config-if)# interface range FastEthernet0/23-24
SW1(config-if)# channel-group 3 mode on

SW2(config)# interface Port-channel1
SW2(config-if)# switchport trunk encapsulation dot1q
SW2(config-if)# switchport mode trunk
SW2(config-if)# interface Port-channel2
```

```

SW2 (config-if) # switchport trunk encapsulation dot1q
SW2 (config-if) # switchport mode trunk
SW2 (config-if) # interface Port-channel3
SW2 (config-if) # switchport trunk encapsulation dot1q
SW2 (config-if) # switchport mode trunk
SW2 (config-if) # interface range FastEthernet0/19-20
SW2 (config-if) # channel-group 1 mode on
SW2 (config-if) # interface range FastEthernet0/21-22
SW2 (config-if) # channel-group 2 mode on
SW2 (config-if) # interface range FastEthernet0/23-24
SW2 (config-if) # channel-group 3 mode on

```

```

SW3 (config) # interface Port-channel1
SW3 (config-if) # switchport trunk encapsulation dot1q
SW3 (config-if) # switchport mode trunk
SW3 (config-if) # interface Port-channel2
SW3 (config-if) # switchport trunk encapsulation dot1q
SW3 (config-if) # switchport mode trunk
SW3 (config-if) # interface range FastEthernet0/19-20
SW3 (config-if) # channel-group 1 mode on
SW3 (config-if) # interface range FastEthernet0/21-22
SW3 (config-if) # channel-group 2 mode on

```

```

SW4 (config) # interface Port-channel1
SW4 (config-if) # switchport trunk encapsulation dot1q
SW4 (config-if) # switchport mode trunk
SW4 (config-if) # interface Port-channel2
SW4 (config-if) # switchport trunk encapsulation dot1q
SW4 (config-if) # switchport mode trunk
SW4 (config-if) # interface range FastEthernet0/19-20
SW4 (config-if) # channel-group 1 mode on
SW4 (config-if) # interface range FastEthernet0/21-22
SW4 (config-if) # channel-group 2 mode on

```

```

SW1# show interfaces port-channel 1 status

```

Port	Name	Status	Vlan	Duplex	Speed	Type
Po1		connected	trunk	a-full	a-100	

SW1# show interfaces port-channel 2 status

Port	Name	Status	Vlan	Duplex	Speed	Type
Po2		connected	trunk	a-full	a-100	

SW1# show interfaces port-channel 3 status

Port	Name	Status	Vlan	Duplex	Speed	Type
Po3		connected	trunk	a-full	a-100	

SW1# show etherchannel summary

Number of channel-groups in use: 3
 Number of aggregators: 3

Group	Port-channel	Protocol	Ports
1	Po1 (SU)	-	Fa0/19 (P) Fa0/20 (P)
2	Po2 (SU)	-	Fa0/21 (P) Fa0/22 (P)
3	Po3 (SU)	-	Fa0/23 (P) Fa0/24 (P)

SW2# show interfaces port-channel 1 status

Port	Name	Status	Vlan	Duplex	Speed	Type
Po1		connected	trunk	a-full	a-100	

SW2# show interfaces port-channel 2 status

Port	Name	Status	Vlan	Duplex	Speed	Type
Po2		connected	trunk	a-full	a-100	

SW2# show interfaces port-channel 3 status

Port	Name	Status	Vlan	Duplex	Speed	Type
Po3		connected	trunk	a-full	a-100	

SW2# show etherchannel summary

Number of channel-groups in use: 3
 Number of aggregators: 3

Group	Port-channel	Protocol	Ports

1	Po1 (SU)	-	Fa0/19 (P)	Fa0/20 (P)
2	Po2 (SU)	-	Fa0/21 (P)	Fa0/22 (P)
3	Po3 (SU)	-	Fa0/23 (P)	Fa0/24 (P)

SW3# **show interface port-channel 1 status**

Port	Name	Status	Vlan	Duplex	Speed	Type
Po1		connected	trunk	a-full	a-100	

SW3# **show interface port-channel 2 status**

Port	Name	Status	Vlan	Duplex	Speed	Type
Po2		connected	trunk	a-full	a-100	

SW3# **show etherchannel summary**

Number of channel-groups in use: 2

Number of aggregators: 2

Group	Port-channel	Protocol	Ports
1	Po1 (SU)	-	Fa0/19 (P) Fa0/20 (P)
2	Po2 (SU)	-	Fa0/21 (P) Fa0/22 (P)

SW4# **show interface port-channel 1 status**

Port	Name	Status	Vlan	Duplex	Speed	Type
Po1		connected	trunk	a-full	a-100	

SW4# **show interface port-channel 2 status**

Port	Name	Status	Vlan	Duplex	Speed	Type
Po2		connected	trunk	a-full	a-100	

SW4# **show etherchannel summary**

Number of channel-groups in use: 2

Number of aggregators: 2

Group	Port-channel	Protocol	Ports
1	Po1 (SU)	-	Fa0/19 (P) Fa0/20 (P)
2	Po2 (SU)	-	Fa0/21 (P) Fa0/22 (P)

- Ensure traffic is distributed on individual Ethernet trunks between switches based on the destination MAC address of individual flows. (2 points)

A common problem with EtherChannels is traffic not being distributed equally among the physical interfaces. Configuring channel load balancing based on the destination MAC address of an individual flow is just one method available to distribute traffic. If you have configured this correctly, as shown in Example 1-5, you have scored 2 points.

EXAMPLE 1-5 Switch 1, 2, 3, and 4 EtherChannel Load Balancing Configuration

```
SW1(config)# port-channel load-balance dst-mac
SW2(config)# port-channel load-balance dst-mac
SW3(config)# port-channel load-balance dst-mac
SW4(config)# port-channel load-balance dst-mac

SW1# show etherchannel load-balance
EtherChannel Load-Balancing Operational State (dst-mac):
Non-IP: Destination MAC address
IPv4: Destination MAC address
IPv6: Destination IP address
```

- Ensure that user interfaces are shut down dynamically by all switches if they toggle excessively; if they remain stable for 35 seconds, they should be reenabled. Configure Fast Ethernet Port 0/10 on each switch so that if multicast traffic is received on this port, the port is automatically disabled. (3 points)

Interfaces that flap can cause problems in a network. Toggling would usually indicate a problem such as a faulty connecting NIC or faulty cable; placing the ports into error disable is a method of stabilizing the environment. To disable a port when multicast traffic is present, you need to configure storm control with the multicast option set to 0. If you have configured this correctly, as shown in Example 1-6, you have scored 3 points.

EXAMPLE 1-6 Switch 1, 2, 3, and 4 Configuration

```
SW1(config)# errdisable recovery cause link-flap
SW1(config)# errdisable recovery interval 35
SW1(config)# interface FastEthernet 0/10
SW1(config-if)# storm-control multicast level 0
SW1(config-if)# storm-control action shutdown
```

```
SW2(config)# errdisable recovery cause link-flap
SW2(config)# errdisable recovery interval 35
SW2(config)# interface FastEthernet 0/10
SW2(config-if)# storm-control multicast level 0
SW2(config-if)# storm-control action shutdown

SW3(config)# errdisable recovery cause link-flap
SW3(config)# errdisable recovery interval 35
SW3(config)# interface FastEthernet 0/10
SW3(config-if)# storm-control multicast level 0
SW3(config-if)# storm-control action shutdown

SW4(config)# errdisable recovery cause link-flap
SW4(config)# errdisable recovery interval 35
SW3(config)# interface FastEthernet 0/10
SW3(config-if)# storm-control multicast level 0
SW3(config-if)# storm-control action shutdown
```

- Fast Ethernet Ports 0/11-17 will be used for future connectivity on each switch. Configure these ports as access ports for VLAN300, which should begin forwarding traffic immediately on connection. Devices connected to these ports will dynamically receive IP addresses from a DHCP server due to be connected to Port 0/18 on sw1. For security purposes this is the only port on the network where DHCP addresses should be allocated from. Ensure the switches intercept the DHCP requests and add the ingress port and VLAN and switch MAC address prior to sending forward to the DHCP server. Limit DHCP requests to 600 packets per minute per user port. (6 points)

This is a DHCP Snooping question. This is a useful security feature that protects the network from rogue DHCP servers. When the DHCP option-82 feature is enabled on the switch with the command **ip dhcp snooping information option**, a subscriber is identified by the switch port through which it connects to the network and by its MAC address. DHCP snooping also facilitates a rate limiting feature for DHCP requests to prevent a DHCP denial of service by excessive false requests from a host, which would have the "gobbler effect" of requesting numerous leases from the same port. The question includes a couple of points that could easily be overlooked if you are suffering from exam pressure, namely the ports are required to be configured with **switchport host** (or by configuring portfast) to set the port mode to access and to forward immediately. The rate limiting is configured in packets per second not per minute as implied, so you would need to pay attention to detail. If you have configured this correctly, as shown in Example 1-7, you have scored 6 points.

EXAMPLE 1-7 Switch 1, 2, 3, and 4 DHCP Snooping Configuration

```
SW1(config)# ip dhcp snooping
SW1(config)# ip dhcp snooping vlan 300
SW1(config)# ip dhcp snooping information option
SW1(config)# int fastEthernet 0/18
SW1(config-if)# ip dhcp snooping trust
SW1(config)# interface range fastEthernet 0/11-17
SW1(config-if-range)# ip dhcp snooping limit rate 10
SW1(config)# interface range fastEthernet 0/11-18
SW1(config-if-range)# switchport host
SW1(config-if-range)# switchport access vlan 300

SW2(config)# ip dhcp snooping
SW2(config)# ip dhcp snooping vlan 300
SW2(config)# ip dhcp snooping information option
SW2(config)# interface range fastEthernet 0/11-17
SW2(config-if-range)# ip dhcp snooping limit rate 10
SW2(config-if-range)# switchport host
SW2(config-if-range)# switchport access vlan 300

SW3(config)# ip dhcp snooping
SW3(config)# ip dhcp snooping vlan 300
SW3(config)# ip dhcp snooping information option
SW3(config)# interface range fastEthernet 0/11-17
SW3(config-if-range)# ip dhcp snooping limit rate 10
SW3(config-if-range)# switchport host
SW3(config-if-range)# switchport access vlan 300

SW4(config)# ip dhcp snooping
SW4(config)# ip dhcp snooping vlan 300
SW4(config)# ip dhcp snooping information option
SW4(config)# interface range fastEthernet 0/11-17
SW4(config-if-range)# ip dhcp snooping limit rate 10
SW4(config-if-range)# switchport host
SW4(config-if-range)# switchport access vlan 300

SW1# sh ip dhcp snooping
```

```

Switch DHCP snooping is enabled
DHCP snooping is configured on following VLANs:
300
Insertion of option 82 is enabled
  circuit-id format: vlan-mod-port
  remote-id format: MAC
Option 82 on untrusted port is not allowed
Verification of hwaddr field is enabled
Interface                Trusted      Rate limit (pps)
-----                -
FastEthernet0/11         no          10
FastEthernet0/12         no          10
FastEthernet0/13         no          10
FastEthernet0/14         no          10
FastEthernet0/15         no          10
FastEthernet0/16         no          10
FastEthernet0/17         no          10
FastEthernet0/18         yes         unlimited

```

- For additional security ensure the user ports on Switches 1-4 and 11-17 can communicate only with the network with IP addresses gained from the DHCP feature configured previously. Use a dynamic feature to ensure the only information forwarded upon connection is DHCP request packets and then any traffic that matches the DHCP IP information received from the DHCP binding for additional security. (3 points)

A complementary feature to DHCP Snooping is IP Source Guard. This feature binds the information received from the DHCP address offered and effectively builds a dynamic VACL on a per port basis to enable only source traffic matched from the DHCP offer to ingress the switch port for additional security. If you have configured this correctly, as shown in Example 1-8, you have scored 3 points.

EXAMPLE 1-8 Switch 1, 2, 3, and 4 IP Source Guard Configuration

```

SW1(config)# interface range fast 0/11-17
SW1(config-if-range)# ip verify source

SW2(config)# interface range fast 0/11-17
SW2(config-if-range)# ip verify source

SW3(config)# interface range fast 0/11-17
SW3(config-if-range)# ip verify source

```

```
SW4(config)# interface range fast 0/11-17
SW4(config-if-range)# ip verify source
```

- R5 and R6 have been preconfigured with IP addresses on their Ethernet interfaces. Configure R4 and its associated switch port accordingly without using secondary addressing to communicate with R5 and R6. Configure R4 with an IP address of 120.100.45.4/24 to communicate with R5, and configure R4 with an IP address of 120.100.46.4/24 to communicate with R6. Configure R4 Gi0/1 and Switch 2 FE0/4 only. (3 points)

This is just a simple trunking question on Switch2 to R4 to enable R4 to connect to VLAN45 and VLAN46. One point to remember is that Switch2 does not have VLAN45 and VLAN46 configured locally within the default configuration, so you will need to create the VLANs locally prior to configuring the trunk. If you have configured this correctly, as shown in Example 1-9, you have scored 3 points.

EXAMPLE 1-9 Switch2 and R4 Trunking Configuration

```
R4(config)# interface GigabitEthernet0/1.45
R4(config-if)# encapsulation dot1Q 45
R4(config-if)# ip address 120.100.45.4 255.255.255.0
R4(config-if)# interface GigabitEthernet0/1.46
R4(config-if)# encapsulation dot1Q 46
R4(config-if)# ip address 120.100.46.4 255.255.255.0

SW2(config)# vlan 45-46
SW2(config)# interface FastEthernet0/4
SW2(config-if)# switchport trunk encapsulation dot1q
SW2(config-if)# switchport trunk allowed vlan 45,46
SW2(config-if)# switchport mode trunk
```

- Your initial configuration has been supplied for the R1-R2-R3 connectivity and R2-R5. Configure each device as per Figure 1-6 to ensure each device is reachable over the Frame Relay network. Only use the indicated DLCIs. (2 points)

The initial Frame Relay configuration has been supplied for you; all you need to add is additional maps on R1 and R2 spokes to enable them to communicate with each other by directing traffic to the Hub router (R3) because the initial configuration uses no inverse arp. Communication between R2 and R5 will work without modification by default. If you have configured this correctly, as shown in Example 1-10, you have scored 2 points.

EXAMPLE 1-10 R1 and R2 Additional Frame Relay Configuration and Testing

```
R1# conf t
R1(config)# int s0/0/0
R1(config-if)# frame-relay map ip 120.100.123.2 103 broadcast

R2# conf t
Enter configuration commands, one per line. End with CNTL/Z.
R2(config)# int s0/0
R2(config-if)# frame-relay map ip 120.100.123.1 203 broadcast

R1# ping 120.100.123.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 120.100.123.2, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 8/8/8 ms
```

Section 2: IPv4 IGP Protocols (22 Points)

Section 2.1: OSPF

- Use a process ID of 1; all OSPF configuration where possible should not be configured under the process ID. Do not change the preconfigured interface types where applicable. The Loopback interfaces of Routers R1, R2, and R3 should be configured to be in Area 0. R4 should be in Area 34 and R5 in Area 5. (2 points)

Recent advances in OSPF have enabled configuration of the network area directly under the interface as opposed to within the OSPF process. Example 1-11 details the OSPF configuration.

EXAMPLE 1-11 OSPF Configuration

```
R1(config)# interface GigabitEthernet 0/1
R1(config-if)# ip ospf 1 area 100
R1(config)# interface Serial 0/0/0
R1(config-if)# ip ospf 1 area 0
R1(config-if)# interface Loopback 0
R1(config-if)# ip ospf 1 area 0
```

```
R2(config)# interface Loopback 0
R2(config-if)# ip ospf 1 area 0
R2(config-if)# interface Serial 0/0
R2(config-if)# ip ospf 1 area 0
R2(config-if)# interface Serial 0/1
R2(config-if)# ip ospf 1 area 5
R2(config-if)# interface FastEthernet 0/1
R2(config-if)# ip ospf 1 area 200

R3(config)# interface loopback 0
R3(config-if)# ip ospf 1 area 0
R3(config-if)# interface Serial 0/0/0
R3(config-if)# ip ospf 1 area 0
R3(config-if)# interface GigabitEthernet 0/0
R3(config-if)# ip ospf 1 area 34

R4(config)# interface Loopback 0
R4(config-if)# ip ospf 1 area 34
R4(config-if)# interface GigabitEthernet 0/0
R4(config-if)# ip ospf 1 area 34
R4(config-if)# interface GigabitEthernet 0/1.45
R4(config-if)# ip ospf 1 area 5

R5(config)# interface Loopback 0
R5(config-if)# ip ospf 1 area 5
R5(config-if)# interface GigabitEthernet 0/0
R5(config-if)# ip ospf 1 area 5
R5(config-if)# interface Serial 0/0/1
R5(config-if)# ip ospf 1 area 5
```

Initial configuration changes the OSPF network interface types on Router R1, R2, and R3 Frame Relay interfaces; this changes the hello and dead interval timers, which results in a mismatch with neighbor relationship never being formed. Example 1-12 shows the differing interface parameters between routers and required configuration on Routers R1 and R3. Because you cannot change the network type, you must manually adjust the OSPF Hello-interval. The most logical place to do this is on the hub Router R3 to ensure a common configuration. If you have configured OSPF correctly, as shown in Examples 1-11 and 1-12, you have scored 2 points.

EXAMPLE 1-12 OSPF Interface Parameters and Configuration

```
R1# show ip ospf interface Serial 0/0/0
Serial0/0/0 is up, line protocol is up
  Internet Address 120.100.123.1/24, Area 0
  Process ID 1, Router ID 120.100.1.1, Network Type POINT_TO_POINT, Cost: 64
  Enabled by interface config, including secondary ip addresses
  Transmit Delay is 1 sec, State POINT_TO_POINT
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    oob-resync timeout 40
    Hello due in 00:00:08
  Supports Link-local Signaling (LLS)
  Cisco NSF helper support enabled
  IETF NSF helper support enabled
  Index 1/2, flood queue length 0
  Next 0x0(0)/0x0(0)
  Last flood scan length is 0, maximum is 0
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 0, Adjacent neighbor count is 0
  Suppress hello for 0 neighbor(s)

R3# show ip ospf interface Serial 0/0/0
Serial0/0/0 is up, line protocol is up
  Internet Address 120.100.123.3/24, Area 0
  Process ID 1, Router ID 120.100.3.1, Network Type POINT_TO_MULTIPOINT, Cost: 64
  Enabled by interface config, including secondary ip addresses
  Transmit Delay is 1 sec, State POINT_TO_MULTIPOINT
  Timer intervals configured, Hello 30, Dead 120, Wait 120, Retransmit 5
    oob-resync timeout 120
    Hello due in 00:00:08
  Supports Link-local Signaling (LLS)
  Cisco NSF helper support enabled
  IETF NSF helper support enabled
  Index 2/2, flood queue length 0
  Next 0x0(0)/0x0(0)
  Last flood scan length is 0, maximum is 0
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 0, Adjacent neighbor count is 0
  Suppress hello for 0 neighbor(s)
```

```
R3# conf t
R3(config)# int Serial 0/0/0
R3(config-if)# ip ospf hello-interval 10
R3# show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
120.100.1.1	0	FULL/ -	00:00:32	120.100.123.1	Serial0/0/0
120.100.2.1	0	FULL/ -	00:00:35	120.100.123.2	Serial0/0/0
120.100.4.1	1	FULL/BDR	00:00:39	120.100.34.4	GigabitEthernet0/0

■ All Loopback networks should not be advertised as host routes. (1 point)

Loopback interfaces within OSPF will by default be advertised as host routes. To manipulate this behavior you need to override the network type that the IOS associates with the Loopback interface. Example 1-13 shows the host routes learned on R2. Note that 120.100.123.3/32 is actually a host route generated by OSPF for the Frame Relay connection, so this is expected behavior and acceptable in the routing table. If you have configured this correctly, as shown in Ex-ample 1-13, you have scored 1 point.

EXAMPLE 1-13 OSPF Loopback Interface Host Routes and Configuration

```
R2# sh ip route | inc /32
```

- O 120.100.5.1/32 [110/65] via 120.100.25.5, 00:04:34, Serial0/1
- O IA 120.100.4.1/32 [110/66] via 120.100.123.3, 00:00:42, Serial0/0
- O 120.100.123.3, 00:01:00, Serial0/0
- O 120.100.1.1/32 [110/129] via 120.100.123.3, 00:01:00, Serial0/0
- O 120.100.3.1/32 [110/65] via 120.100.123.3, 00:01:00, Serial0/0
- O 120.100.123.3/32 [110/64] via 120.100.123.3, 00:01:00, Serial0/0

```
R1# conf t
R1(config)# int Loopback 0
R1(config-if)# ip ospf network point-to-point
R2#conf t
R2(config)# interface Loopback 0
R2(config-if)# ip ospf network point-to-point
R3# conf t
R3(config)# int Loopback 0
```

```
R3(config-if)# ip ospf network point-to-point
```

```
R4# conf t
```

```
R4(config)# int Loopback 0
```

```
R4(config-if)# ip ospf network point-to-point
```

```
R5# conf t
```

```
R4(config)# int Loopback 0
```

```
R4(config-if)# ip ospf network point-to-point
```

```
R2# sh ip route ospf 1 | include /24
```

```
150.100.0.0/24 is subnetted, 2 subnets
O IA 120.100.4.0/24 [110/66] via 120.100.123.3, 00:00:43, Serial0/0
O 120.100.5.0/24 [110/65] via 120.100.25.5, 00:01:40, Serial0/1
O 120.100.1.0/24 [110/129] via 120.100.123.3, 00:00:43, Serial0/0
O 120.100.3.0/24 [110/65] via 120.100.123.3, 00:00:43, Serial0/0
O 120.100.45.0/24 [110/65] via 120.100.25.5, 00:01:40, Serial0/1
O IA 120.100.34.0/24 [110/65] via 120.100.123.3, 00:00:43, Serial0/0
O IA 120.100.100.0/24 [110/129] via 120.100.123.3, 00:00:09, Serial0/0
```

- Ensure that R1 does not advertise the preconfigured secondary address under interface Gigabit 0/1 of 120.100.100.1/24 to the OSPF network. Do not use any filtering techniques to achieve this. (2 points)

The associated behavior with configuring OSPF directly under the interface is that it will by default advertise any secondary addresses assigned to the interface. R1 has a preconfigured secondary address on interface Gigabit 0/1 that is therefore advertised. Because you cannot filter this advertisement, you need to inform OSPF not to include the secondary addresses under the interface command. If you have configured this correctly, as shown in Example 1-14, you have scored 2 points.

EXAMPLE 1-14 OSPF Secondary Address Advertisement and Configuration

```
R1# show ip ospf int GigabitEthernet 0/1
```

```
GigabitEthernet0/1 is up, line protocol is up
Internet Address 150.100.1.1/24, Area 100
Process ID 1, Router ID 120.100.1.1, Network Type BROADCAST, Cost: 1
Enabled by interface config, including secondary ip addresses
Transmit Delay is 1 sec, State DR, Priority 1
Designated Router (ID) 120.100.1.1, Interface address 150.100.1.1
```

```
No backup designated router on this network
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit
  5 oob-resync timeout 40
  Hello due in 00:00:00
Supports Link-local Signaling (LLS)
Cisco NSF helper support enabled
IETF NSF helper support enabled
Index 1/1, flood queue length 0
Next 0x0(0)/0x0(0)
Last flood scan length is 0, maximum is 0
Last flood scan time is 0 msec, maximum is 0 msec
Neighbor Count is 0, Adjacent neighbor count is 0
Suppress hello for 0 neighbor(s)
```

```
R1(config)# interface GigabitEthernet 0/1
R1(config-if)# ip ospf 1 area 100 secondaries none
```

```
R2# sh ip route 120.100.100.0
% Subnet not in table
```

- R5 should use the Frame Relay link within Area 5 for its primary communication to the OSPF network. If this network should fail either at Layer 1 or Layer 2, R5 should form a neighbor relationship with R4 under Area 5 to maintain connectivity. Your solution should be dynamic, ensuring that while the Area 5 Frame Relay link is operational, there is no neighbor relationship between R4 and R5; however, the Ethernet interfaces of R4 and R5 must remain up. To confirm the operational status of the Frame Relay network, you should ensure that the serial interface of R5 is reachable by configuration of R5. You are permitted to define neighbor statements between R5 and R4. (4 points)

This is a complex scenario that can consume your time, but all the clues are in the question, so some lateral thinking is required. You can rule out a backup interface solution because the Ethernet needs to remain up, and the solution must cater for Layer 1 and Layer 2 rather than purely Layer 1. Similarly, a demand scenario is also out because this would involve a neighbor relationship being formed. You are also requested to confirm operational status of the Frame Relay interface on R5 with your overall solution being dynamic. This would take a great deal of effort and trial and error, but you will find that you can use the IP SLA feature to monitor the IP address of the Frame Relay interface on R5 by R5 itself. If this responds to the automatic polling with ICMP, you know the frame relay is up at Layers 1 and 2. (Layer 2 would also need to be up for a valid response because the ICMP packet would be sent over the Frame Relay network, and a local map to R5's own IP address is required for this.) If the polling fails, you know the interface is down. IP SLA

can then be used to inform the router, and a forwarding decision can be manipulated; this feature is known as Policy-Based Routing (PBR) support with multiple Tracking Options. This gives PBR access to all the objects that are available through the tracking process.

The tracking process provides the ability to track individual objects, such as ICMP ping reachability, and inform the required PBR process when an object state changes. In summary, if the object status changes, R5 can simply manipulate the way it sends traffic by policy routing. The traffic it manipulates needs to be OSPF that should be directed to R4 to form the adjacency over the Ethernet network (VLAN45), so when R5 Frame Relay is up and running, we just need to break the adjacency between R5 and R4. When the Frame Relay fails, we need to allow the adjacency between R5 and R4 to form. The first step in this solution is to configure the IP SLA object tracking on R5. Remember the additional map is needed locally, so it can ping its own serial interface; this configuration is detailed in Example 1-15.

EXAMPLE 1-15 R5 IP SLA Configuration and Status

```
R5(config)# interface Serial0/0/1
R5(config-if)# frame-relay map ip 120.100.25.5 512 broadcast
R5(config-if)# exit
R5(config)# ip sla 1
R5(config-ip-sla)# icmp-echo 120.100.25.5
R5(config-ip-sla-echo)# ip sla schedule 1 life forever start-time now
R5(config)# track 1 rtr 1 reachability
```

```
R5# show ip sla statistics
```

```
Round Trip Time (RTT) for      Index 1
  Latest RTT: 4 milliseconds
Latest operation start time: *21:17:10.683 UTC Mon Feb 19 2007
Latest operation return code: OK
Number of successes: 2
Number of failures: 0
Operation time to live: Forever
```

NOTE

OSPF should have already been configured between R4 and R5 within your original peering configuration. The neighbor adjacency takes a while waiting for the dead time to expire (120 seconds after changing of the OSPF network type).

OSPF needs to be configured between R4 and R5 with manual neighbor statements as directed in the question, which ensures the routers unicast traffic to each other. To do this you need to change the network type to nonbroadcast. The unicast traffic between neighbors can be identified by an ACL that the PBR process can match, and then instead of allowing normal traffic flow between R5 and R4 to form the neighbor relationship, the next hop can be modified and as the OSPF TTL is set to 1 by default, the traffic will effectively be dropped by the next hop and the OSPF between R5

and R4 will never establish. Similarly, when the object tracking fails, the PBR process will be overridden and traffic can flow as normal. This will then allow R5 and R4 to form an OSPF adjacency. So by using the PBR command **set ip next-hop verify-availability 120.100.25.2 10 track 1**, R5 can forward normal OSPF traffic to 120.100.25.2 (R2 Frame Relay to effectively discard the traffic) if the tracked object (1) is up. If the object status changes to down, the PBR process is informed, and the OPSF traffic to 120.100.25.2 would follow the usual next hop. R5 must be configured to locally policy route traffic because normal PBR behavior is for traffic manipulation for traffic that flows through the router rather than traffic generated by the router itself. Example 1-16 shows the required OSPF configuration on R4 and R5, the PBR on R5, a debug of R2 sending TTL expired to R5 after the OSPF traffic is sent to R2 instead of R5, and the resulting neighbor partial adjacency that is formed between R4 and R5.

EXAMPLE 1-16 R4 and R5 OSPF and PBR Configuration

```
R4(config)# interface GigabitEthernet0/1.45
R4(config-if)# ip ospf network non-broadcast
R4(config-if)# router ospf 1
R4(config-router)# neighbor 120.100.45.5

R5(config)# interface GigabitEthernet0/0
R5(config-if)# ip ospf network non-broadcast
R5(config-if)# router ospf 1
R5(config-router)# neighbor 120.100.45.4
R5(config-router)# exit
R5(config)# access-list 100 permit ospf host 120.100.45.5 host 120.100.45.4
R5(config)# route-map TEST permit 10
R5(config-route-map)# match ip address 100
R5(config-route-map)# set ip next-hop verify-availability 120.100.25.2 10 track 1
R5(config-route-map)# interface GigabitEthernet0/0
R5(config-if)# ip policy route-map TEST
R5(config-if)# exit
R5(config)# ip local policy route-map TEST

R2# debug ip icmp
ICMP packet debugging is on
R2#
*Feb 26 22:17:12.847: ICMP: time exceeded (time to live) sent to 120.100.45.5 (d
```

```
est was 120.100.45.4)
```

```
R2#
```

```
R5# show ip ospf neigh
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
120.100.2.1	0	FULL/ -	00:00:37	120.100.25.2	Serial0/0/1
120.100.4.1	1	INIT/DROTHER	00:01:45	120.100.45.4	GigabitEthernet0/0

Example 1-17 shows the OSPF adjacency formed when the Frame Relay between R2 and R5 is shut down on R5. The PBR is overridden and normal routing occurs because the next hop is not verified by the object tracking. Your routing table needs to be an exact replica as that shown in Example 1-17. You must remember that when an OSPF adjacency forms between R5 and R2, you are joining Area 5 into Area 34 and a virtual-link between R3 and R4 is required to extend area 0. If you hadn't configured a virtual-link it would have been an easy mistake that would take your points away. A difficult question but a good one to practice with and examine how features operate and interact with each other, you may have been scratching your head or cursing me but I'd be surprised if you didn't learn something new from this question.

If you configured this correctly, including the virtual link, you have scored 4 points—definitely a question worth leaving to the end of your exam when hopefully you have time left over to experiment.

EXAMPLE 1-17 R3 and R4 OSPF Virtual Link Configuration and R5 Test

```
R3(config)# router ospf 1
R3(config-router)# area 34 virtual-link 120.100.4.1

R4(config)# router ospf 1
R4(config-router)# area 34 virtual-link 120.100.3.1

R5(config)# interface s0/0/1
R5(config-if)# shut
R5(config-if)#
*Jan  2 21:58:16.811: %OSPF-5-ADJCHG: Process 1, Nbr 120.100.2.1 on Serial0/0/1 from FULL to DOWN, Neighbor
Down: Interface down or detached
*Jan  2 21:58:18.807: %LINK-5-CHANGED: Interface Serial0/0/1, changed state to administratively down
*Jan  2 21:58:19.807: %LINEPROTO-5-UPDOWN: Line protocol on Interface Serial0/0/1, changed state to down
R5(config-if)# do show ip ospf neigh
```

```

Neighbor ID      Pri   State           Dead Time           Address             Interface
N/A              0    ATTEMPT/DROTHER 00:00:33           120.100.45.4      GigabitEthernet0/0
R5(config-if)#
*Jan  2 21:59:43.547: %OSPF-5-ADJCHG: Process 1, Nbr 0.0.0.0 on GigabitEthernet0/0 from ATTEMPT to
DOWN, Neighbor Down: Dead timer expired
R5(config-if)#
*Jan  2 22:00:08.135: %OSPF-5-ADJCHG: Process 1, Nbr 120.100.4.1 on GigabitEthernet0/0 from LOADING to FULL,
Loading Done
R5(config-if)#

R5# sh ip route ospf
      150.100.0.0/24 is subnetted, 3 subnets
O IA   150.100.2.0 [110/67] via 120.100.45.4, 00:09:41, GigabitEthernet0/0
O IA   150.100.1.0 [110/67] via 120.100.45.4, 00:09:41, GigabitEthernet0/0
      120.0.0.0/8 is variably subnetted, 13 subnets, 2 masks
O IA   120.100.25.0/24
          [110/130] via 120.100.45.4, 00:09:41, GigabitEthernet0/0
O IA   120.100.4.1/32 [110/2] via 120.100.45.4, 00:09:41, GigabitEthernet0/0
O IA   120.100.1.0/24 [110/67] via 120.100.45.4, 00:09:41, GigabitEthernet0/0
O IA   120.100.2.0/24 [110/67] via 120.100.45.4, 00:09:41, GigabitEthernet0/0
O IA   120.100.3.0/24 [110/3] via 120.100.45.4, 00:09:41, GigabitEthernet0/0
O IA   120.100.34.0/24 [110/2] via 120.100.45.4, 00:09:41, GigabitEthernet0/0
O IA   120.100.123.3/32
          [110/2] via 120.100.45.4, 00:09:41, GigabitEthernet0/0
O IA   120.100.123.0/24
          [110/130] via 120.100.45.4, 00:09:41, GigabitEthernet0/0

```

Section 2.2: EIGRP

- Configure EIGRP using an AS number of 1. The Loopback interfaces of all routers and switches should be advertised within EIGRP. (2 points)

Not a difficult question by any means; just one that has a magnitude of configuration and sets up your EIGRP network for the following questions. You need to remember to include your preconfigured Loopback interfaces and enable routing on the Layer 3 switches. Use the **show ip eigrp neighbor** command to verify your peering prior to moving onto the next question. If you have configured this correctly, as shown in Example 1-18, you have scored 2 points.

EXAMPLE 1-18 EIGRP Configuration

```
R4# sh run | beg eigrp
router eigrp 1
  network 120.100.4.1 0.0.0.0
  network 120.100.45.4 0.0.0.0
  network 120.100.46.4 0.0.0.0
  no auto-summary
```

```
R5# sh run | beg eigrp
router eigrp 1
  passive-interface Loopback0
  network 120.100.5.1 0.0.0.0
  network 120.100.45.5 0.0.0.0
  network 150.100.3.5 0.0.0.0
  no auto-summary
```

```
R6# sh run | beg eigrp
router eigrp 1
  network 120.100.6.1 0.0.0.0
  network 120.100.46.6 0.0.0.0
  network 150.100.3.6 0.0.0.0
  no auto-summary
```

```
SW1(config)# ip routing
SW1(config)# exit
SW1# sh run | beg eigrp
router eigrp 1
  network 120.100.7.1 0.0.0.0
  network 150.100.3.7 0.0.0.0
  no auto-summary
```

```
SW2(config)# ip routing
SW2(config)# exit
SW2# sh run | beg eigrp
router eigrp 1
  network 120.100.8.1 0.0.0.0
```

```
network 150.100.3.8 0.0.0.0
no auto-summary
```

```
SW3(config)# ip routing
```

```
SW3(config)# exit
```

```
SW3# sh run | beg eigrp
```

```
router eigrp 1
network 120.100.9.1 0.0.0.0
network 150.100.3.9 0.0.0.0
no auto-summary
```

```
SW4(config)#ip routing
```

```
SW4(config)#exit
```

```
SW4# sh run | beg eigrp
```

```
router eigrp 1
network 120.100.10.1 0.0.0.0
network 150.100.3.10 0.0.0.0
no auto-summary
```

-
- Ensure that R4 does not install any of the EIGRP Loopback routes from any of the switches into its routing table; as such, these routes should also not be present in the OSPF network post redistribution. Do not use any route-filtering ACLs, prefix lists, or admin distance manipulation to achieve this, and perform configuration only on R4. (3 points)

A distribute or prefix list would have been the obvious choice here but this is not permitted. Upon close inspection of the Loopback routes within Example 1-19, you will notice that the routes have a hop count of 2 associated with them. Hop count isn't something you would naturally assimilate with EIGRP, but you can configure the process to ignore routes received with a hop count larger than a configured threshold with the command **metric maximum-hops**. By configuring the maximum hop count of 1 on R4, you can simply stop the Loopback routes from entering the process. If you have configured this correctly, as shown in Example 1-19, you have scored 3 points.

EXAMPLE 1-19 EIGRP maximum-hops Configuration

```
R4# show ip route eigrp
150.100.0.0/24 is subnetted, 3 subnets
D    150.100.3.0
        [90/30720] via 120.100.46.6, 00:00:10, GigabitEthernet0/1.46
        [90/30720] via 120.100.45.5, 00:00:10, GigabitEthernet0/1.45
120.0.0.0/8 is variably subnetted, 16 subnets, 2 masks
```

```
D      120.100.8.0/24
      [90/158720] via 120.100.46.6, 00:00:10, GigabitEthernet0/1.46
      [90/158720] via 120.100.45.5, 00:00:10, GigabitEthernet0/1.45
D      120.100.9.0/24
      [90/158720] via 120.100.46.6, 00:00:10, GigabitEthernet0/1.46
      [90/158720] via 120.100.45.5, 00:00:10, GigabitEthernet0/1.45
D      120.100.10.0/24
      [90/158720] via 120.100.46.6, 00:01:07, GigabitEthernet0/1.46
      [90/158720] via 120.100.45.5, 00:01:07, GigabitEthernet0/1.45
D      120.100.5.0/24
      [90/156160] via 120.100.45.5, 00:00:10, GigabitEthernet0/1.45
D      120.100.6.0/24
      [90/156160] via 120.100.46.6, 00:00:10, GigabitEthernet0/1.46
D      120.100.7.0/24
      [90/158720] via 120.100.46.6, 00:00:10, GigabitEthernet0/1.46
      [90/158720] via 120.100.45.5, 00:00:10, GigabitEthernet0/1.45
```

R4# show ip route 120.100.8.0

```
Routing entry for 120.100.8.0/24
  Known via "eigrp 1", distance 90, metric 158720, type internal
  Redistributing via ospf 1, eigrp 1
  Advertised by ospf 1 metric 5000 subnets
  Last update from 120.100.46.6 on GigabitEthernet0/1.46, 00:00:15 ago
  Routing Descriptor Blocks:
  * 120.100.46.6, from 120.100.46.6, 00:00:15 ago, via GigabitEthernet0/1.46
    Route metric is 158720, traffic share count is 1
    Total delay is 5200 microseconds, minimum bandwidth is 100000 Kbit
    Reliability 255/255, minimum MTU 1500 bytes
    Loading 1/255, Hops 2
```

R4# show ip route 120.100.9.0

```
Routing entry for 120.100.9.0/24
  Known via "eigrp 1", distance 90, metric 158720, type internal
  Redistributing via ospf 1, eigrp 1
  Advertised by ospf 1 metric 5000 subnets
  Last update from 120.100.46.6 on GigabitEthernet0/1.46, 00:00:25 ago
  Routing Descriptor Blocks:
  * 120.100.46.6, from 120.100.46.6, 00:00:25 ago, via GigabitEthernet0/1.46
    Route metric is 158720, traffic share count is 1
```

```
Total delay is 5200 microseconds, minimum bandwidth is 100000 Kbit
Reliability 255/255, minimum MTU 1500 bytes
Loading 1/255, Hops 2
```

```
R4(config)# router eigrp 1
R4(config-router)# metric maximum-hops 1
R4(config-router)# do show ip route eigrp
 150.100.0.0/24 is subnetted, 3 subnets
D    150.100.3.0
        [90/30720] via 120.100.46.6, 00:00:04, GigabitEthernet0/1.46
        [90/30720] via 120.100.45.5, 00:00:04, GigabitEthernet0/1.45
 120.0.0.0/8 is variably subnetted, 13 subnets, 2 masks
D    120.100.5.0/24
        [90/156160] via 120.100.45.5, 00:00:04, GigabitEthernet0/1.45
D    120.100.6.0/24
        [90/156160] via 120.100.46.6, 00:00:04, GigabitEthernet0/1.46
```

- R4 will have dual equal cost routes to VLAN300 (network 150.100.3.0) from R5 and R6. Ensure R4 sends traffic to this destination network to R5 rather than load sharing; should the route from R5 become unavailable, traffic should be sent to R6. You may not policy route, alter the bandwidth, or delay statements on R4's interfaces or use an offset list. Perform your configuration on R4 only. Your solution should be applied to all routes received from R5 and R6 as opposed to solely the route to network VLAN300. (3 points)

To receive identical routes your topology must have identical interface types or bandwidth statements used on R4, R5, and R6. Example 1-20 shows the VLAN300 route (150.100.3.0/24) received on R4 from both R5 and R6 with a metric of 30720. If you wanted to manipulate this route the usual best practice method would be to modify the bandwidth or delay on one of the Ethernet interfaces, but this is not permitted. In fact, you are only left with one method that can be applied on R4, which will influence all routes from R5 and R6, as opposed to just this individual route. A route-map is required to override the EIGRP assigned metrics assigned to routes on one interface by manipulating the bandwidth assigned to Gigabit 1/0.45. Gigabit 1/0.46 will, by default, have a lower bandwidth assigned to routes received from it from the permit 20 statement in the route-map. The route-map is applied inbound to the process as a distribute-list. Example 1-20 also shows that when the interface Gigabit 0/0 is shut down on R5 that the route for VLAN300 is still received from R6 (R4's feasible successor), so the route is still available but with a different metric. If you have configured this correctly, as shown in Example 1-20, you have scored 3 points. (You could have also manipulated the delay within the route-map or created a statement for each individual interface as opposed to just Gigabit 1/0.45.)

EXAMPLE 1-20 EIGRP Metric Manipulation Configuration

```
R4# sh ip route 150.100.3.0
Routing entry for 150.100.3.0/24
  Known via "eigrp 1", distance 90, metric 30720, type internal
  Redistributing via ospf 1, eigrp 1
  Advertised by ospf 1 metric 5000 subnets
  Last update from 120.100.45.5 on GigabitEthernet0/1.45, 00:25:40 ago
  Routing Descriptor Blocks:
  * 120.100.46.6, from 120.100.46.6, 00:25:40 ago, via GigabitEthernet0/1.46
    Route metric is 30720, traffic share count is 1
    Total delay is 200 microseconds, minimum bandwidth is 100000 Kbit
    Reliability 254/255, minimum MTU 1500 bytes
    Loading 1/255, Hops 1
  120.100.45.5, from 120.100.45.5, 00:25:40 ago, via GigabitEthernet0/1.45
    Route metric is 30720, traffic share count is 1
    Total delay is 200 microseconds, minimum bandwidth is 100000 Kbit
    Reliability 252/255, minimum MTU 1500 bytes
    Loading 1/255, Hops 1
```

```
R4(config)# route-map CHANGEMETRIC permit 10
R4(config-route-map)# match interface gigabitEthernet 0/1.45
R4(config-route-map)# set metric 2000 10 255 1 1500
R4(config-route-map)# route-map CHANGEMETRIC permit 20
R4(config-route-map)# set metric 1000 10 255 1 1500
R4(config-route-map)# router eigrp 1
R4(config-router)# distribute-list route-map CHANGEMETRIC in
R4(config-router)# ^Z
R4# clear ip route *
R4# sh ip route 150.100.3.0
Routing entry for 150.100.3.0/24
  Known via "eigrp 1", distance 90, metric 1282560, type internal
  Redistributing via ospf 1, eigrp 1
  Advertised by ospf 1 metric 5000 subnets
  Last update from 120.100.45.5 on GigabitEthernet0/1.45, 00:03:10 ago
  Routing Descriptor Blocks:
  * 120.100.45.5, from 120.100.45.5, 00:03:10 ago, via GigabitEthernet0/1.45
    Route metric is 1282560, traffic share count is 1
    Total delay is 100 microseconds, minimum bandwidth is 2000 Kbit
```

```
Reliability 255/255, minimum MTU 1500 bytes
Loading 1/255, Hops 1
```

```
R5(config)# int gig0/0
R5(config-if)# shutdown
```

```
R4# sh ip route 150.100.3.0
Routing entry for 150.100.3.0/24
  Known via "eigrp 1", distance 90, metric 2562560, type internal
  Redistributing via ospf 1, eigrp 1
  Advertised by ospf 1 metric 5000 subnets
  Last update from 120.100.46.6 on GigabitEthernet0/1.46, 00:00:10 ago
  Routing Descriptor Blocks:
  * 120.100.46.6, from 120.100.46.6, 00:00:10 ago, via GigabitEthernet0/1.46
    Route metric is 2562560, traffic share count is 1
    Total delay is 100 microseconds, minimum bandwidth is 1000 Kbit
    Reliability 255/255, minimum MTU 1500 bytes
    Loading 1/255, Hops 1
```

Section 2.3: Redistribution

- Perform mutual redistribution of IGP protocols on R4. All routes should be accessible with the exception of the switch Loopback networks because these should not be visible via R4 from an earlier question. EIGRP routes redistributed within the OSPF network should remain with a fixed cost of 5000 throughout the network. (3 points)

A simple redistribution question for the warm-up lab, you have only a single redistribution point (R4), so have no concerns when using protocols such as EIGRP and OSPF, with their inherent protection against routing loops. The fixed cost of 5000 is achieved by advertising redistributed routes into OSPF using a metric-type of 2, which is the default, so no specific configuration is required for this. The only points you need to consider when redistributing into OSPF are to use the **subnets** command to ensure classless redistribution and to use default-metrics in each protocol. If you have configured this correctly, as shown in Example 1-21, you have scored 3 points.

EXAMPLE 1-21 R4 Redistribution Configuration and Verification

```
R4(config)# router eigrp 1
R4(config-router)# redistribute ospf 1
R4(config-router)# default-metric 10000 100 255 1 1500
R4(config-router)# router ospf 1
R4(config-router)# redistribute eigrp 1 subnets
```

```

R4(config-router)# default-metric 5000

R1# show ip route ospf | include E2
O E2 150.100.3.0 [110/5000] via 120.100.123.3, 00:00:46, Serial10/0/0
O E2 120.100.6.0/24 [110/5000] via 120.100.123.3, 00:00:46, Serial10/0/0
O E2 120.100.46.0/24 [110/5000] via 120.100.123.3, 00:00:46, Serial10/0/0

SW1# show ip route eigrp | include EX
D EX 150.100.2.0 [170/284416] via 150.100.3.6, 00:01:43, Vlan300
D EX 150.100.1.0 [170/284416] via 150.100.3.6, 00:01:43, Vlan300
D EX 120.100.25.0/24 [170/284416] via 150.100.3.6, 00:01:43, Vlan300
D EX 120.100.1.0/24 [170/284416] via 150.100.3.6, 00:01:43, Vlan300
D EX 120.100.2.0/24 [170/284416] via 150.100.3.6, 00:01:43, Vlan300
D EX 120.100.3.0/24 [170/284416] via 150.100.3.6, 00:01:43, Vlan300
D EX 120.100.34.0/24 [170/284416] via 150.100.3.6, 00:01:43, Vlan300
D EX 120.100.123.3/32 [170/284416] via 150.100.3.6, 00:01:43, Vlan300
D EX 120.100.123.0/24 [170/284416] via 150.100.3.6, 00:01:44, Vlan300

```

- Configure R4 to only redistribute up to five EIGRP routes, and generate a system warning when the fourth route is redistributed. Do not use any access-lists in your solution. (2 points).

You can limit the number of prefixes redistributed into OSPF and generate a warning when the number of prefixes reaches a defined maximum by use of the **redistribute maximum-prefix** command. To generate the warning on the fourth route, you must configure a percentage threshold (80 percent). If you have configured this correctly, as shown in Example 1-22, you have scored 2 points.

EXAMPLE 1-22 R4 Prefix Configuration

```

R4(config)# router ospf 1
R4(config-router)# redistribute maximum-prefix 5 80

```

Section 3: BGP (14 Points)

- Configure iBGP peering as follows: R1-R3, R2-R3, R6-R5, Sw1-R6, and Sw1-R5. Use minimal configuration and use Loopback interfaces for your peering. Configure eBGP peering as follows: R3-R4, R4-R6, R4-R5, and R5-R2. Use minimal configuration and use Loopback interfaces for your peering with the exception of R4 to R5. (2 points) Use the AS numbers supplied in Figure 1-9. For your eBGP peering on R3, use the TTL security fea-

ture, which will not permit a session from R4 to become established if R4 is more than 2 hops away. This feature must be configured only on R3 and not on R4. (2 points)

Easy peering points to begin with but lots of typing to earn them. You must remember to use peer groups to minimize configuration where possible, namely on R3, R6, and Switch1, and follow the peering instructions closely as these are relevant for the following questions. You should have noticed that R3 was required to be a route reflector for iBGP peers R1 and R2 in AS10 and that **no synchronization** is required because the underlying IGP is not redistributed into BGP. Remember to verify your peering with the **show ip bgp neighbor** command. The peering becomes complicated when the TTL security feature is enabled by use of the command **neighbor 120.100.4.1 ttl-security hops 2** on R3. This command is a neat feature that will not permit the peering session if the received neighbor TTL value is less than 253 in this case, which would suggest that the incoming session could be some form of remote attack with spoofed source IP address of the original neighbor. Because you are not permitted to configure the same feature on R4, the peering will of course break, even if you have configured the ebgp multihop feature on R4 with a value of 2. (Of course this will simply increment the TTL value from a default value of 0.)

Example 1-23 shows a debug on R3 for the ebgp peering; the field highlighted is the TTL Hex value displayed from the hidden command (dump) when performing the debug. You need to get the Hex value to FD (253 decimal) to show R3 that the R4 can only be a maximum of two hops away by configuring the multihop value to 255 on R4. If you have configured this correctly, as shown in Example 1-23, you have scored 2 points.

EXAMPLE 1-23 BGP Peering Configuration

```
R1# sh run | begin bgp
router bgp 10
  no synchronization
  neighbor 120.100.3.1 remote-as 10
  neighbor 120.100.3.1 update-source Loopback0
  no auto-summary

R2# sh run | begin bgp
router bgp 10
  no synchronization
  neighbor 120.100.3.1 remote-as 10
  neighbor 120.100.5.1 remote-as 300
  neighbor 120.100.5.1 ebgp-multihop 2
  neighbor 120.100.5.1 update-source Loopback0
  no auto-summary
```

```
R3# sh run | begin bgp
router bgp 10
  no synchronization
  neighbor IBGP peer-group
  neighbor IBGP remote-as 10
  neighbor IBGP update-source Loopback0
  neighbor IBGP route-reflector-client
  neighbor 120.100.1.1 peer-group IBGP
  neighbor 120.100.2.1 peer-group IBGP
  neighbor 120.100.4.1 remote-as 200
  neighbor 120.100.4.1 ttl-security hops 2
  neighbor 120.100.4.1 update-source Loopback0
  no auto-summary
```

```
R4# sh run | begin bgp
router bgp 200
  no synchronization
  neighbor 120.100.3.1 remote-as 10
  neighbor 120.100.3.1 ebgp-multihop 2
  neighbor 120.100.3.1 update-source Loopback0
  neighbor 120.100.6.1 remote-as 300
  neighbor 120.100.6.1 ebgp-multihop 2
  neighbor 120.100.6.1 update-source Loopback0
  neighbor 120.100.45.5 remote-as 300
  no auto-summary
```

```
R3(config)# access-list 100 permit ip host 120.100.4.1 host 120.100.3.1
```

```
R3(config)# exit
```

```
R3# debug ip packet 100 detail dump
```

```
IP packet debugging is on (detailed) (dump) for access list 100
```

```
R3# TCP src=42692, dst=179, seq=2600279946, ack=0, win=163
84 SYN
0F400C00:                C204 07400000                B..@..
0F400C10: C20211E0 00100800 45C0002C 6A870000  B..`....E@.,j...
0F400C20: 0106467E 01010101 03030303 A6C400B3  ..F~.....&D.3
0F400C30: 9AFD1F8A 00000000 60024000 F1BB0000  .}.....`.@.q;..
0F400C40: 02040218                ....
```

```
! The TTL from R4 is decremented to 01 Hex = 01 decimal as R4 has ebgp-multihop 2
```

```
! configured and the BGP session will not be established as R3 has the TTL security
! check enabled, from R3's perspective R4 could be 254 hops away!
! Configure R4 so the TTL value will read 253 decimal (FD hex) by configuring an
! ebgp multihop value of 255 (this value will decrement down to 253 when it is
! processed by R3).
```

```
R4(config)# router bgp 200
R4(config)# neighbor 120.100.3.1 ebgp-multihop 255
```

```
R3# TCP src=44109, dst=179, seq=3825370469, ack=3209854606
, win=16263 ACK
0F7CBB60: C204 07400000 B..@..
0F7CBB70: C20211E0 00100800 45C00028 8C9A0000 B..`....E@.(....
0F7CBB80: FD06286E 01010101 03030303 AC4D00B3 }. (n.....,M.3
0F7CBB90: E4028565 BF527E8E 50103F87 13FC0000 d..e?R~.P.?..|..
0F7CBBA0:
```

```
! Now a hex value of FD (253 Decimal) can be seen at R3 from R4, this shows that R4
! can not be further than 2 hops away from R3 and the security check passes and BGP
! is established.
```

```
R3# sh ip bgp neighbor | include hops | TTL
External BGP neighbor may be up to 2 hops away.
Connection is ECN Disabled, Minimum incoming TTL 253, Outgoing TTL 255
```

```
R5# sh run | begin bgp
router bgp 300
no synchronization
neighbor 120.100.2.1 remote-as 10
neighbor 120.100.2.1 ebgp-multihop 2
neighbor 120.100.2.1 update-source Loopback0
neighbor 120.100.6.1 remote-as 300
neighbor 120.100.6.1 update-source Loopback0
neighbor 120.100.45.6 remote-as 200
neighbor 150.100.3.7 remote-as 300
no auto-summary
```

```
R6# sh run | beg bgp
router bgp 300
```

```
no synchronization neighbor
IBGP peer-group neighbor
IBGP remote-as 300
neighbor IBGP update-source Loopback0
neighbor 120.100.4.1 remote-as 200
neighbor 120.100.4.1 ebgp-multihop 2
neighbor 120.100.4.1 update-source Loopback0
neighbor 120.100.5.1 peer-group IBGP
neighbor 150.100.3.7 peer-group IBGP
no auto-summary
```

```
SW1# sh run | begin bgp
router bgp 300
no synchronization neighbor
IBGP peer-group neighbor
IBGP remote-as 300
neighbor 120.100.5.1 peer-group IBGP
neighbor 120.100.6.1 peer-group IBGP
no auto-summary
```

- AS200 is to be used as a backup transit network for traffic between AS100 and AS300; as such if the FR network between R5 and R2 fails, ensure the peering between R2 and R5 is not maintained via the Ethernet network. Do not use any ACL type restrictions or change the existing peering. (2 points)

As R2 and R5 peer to each other using their Loopback interfaces, the peering is maintained if the Frame Relay network between R2 and R5 fails. Example 1-24 shows the path taken between R5 and R2 when the Frame Relay interface is shut down on R5. To break the peering without using ACLs, you simply need to ensure the **ebgp-multihop** count used in the original peering is set at 2 and no greater. Example 1-24 also shows the ICMP debug with the TTL expiration messages, which indicate the peering will have failed, even though there is IP connectivity between Loopbacks. If your ebgp-multihop count is set at 2 between R2 and R5, you have scored 2 points.

EXAMPLE 1-24 eBGP TTL Expiration

```
R5(config) #int s0/0/1
R5(config-if)# shut

R5# trace 120.100.2.1

Type escape sequence to abort.
```

```
Tracing the route to 120.100.2.1

  1 120.100.45.4 0 msec 0 msec 0 msec
  2 120.100.34.3 0 msec 4 msec 0 msec
  3 120.100.123.2 4 msec * 4 msec
```

```
R5# debug ip icmp
ICMP packet debugging is on
R5#
*Jan 17 21:32:32.455: ICMP: time exceeded rcvd from 120.100.34.3
R5#
*Jan 17 21:32:34.179: ICMP: time exceeded rcvd from 120.100.34.3
R5#
```

```
R2# debug ip icmp
ICMP packet debugging is on
R2#
Jan 17 21:26:11.310: ICMP: time exceeded rcvd from 120.100.34.4
R2#
Jan 17 21:26:13.306: ICMP: time exceeded rcvd from 120.100.34.4
```

- Configure a new Loopback interface 2 on R2 of 130.100.200.1/24, and advertise this into BGP using the network command. Configure R2 in such a way that if the Frame Relay connection between R2 and R5 fails, AS300 no longer receives this route. Do not use any filtering between neighbors to achieve this or neighbor-specific commands. (3 points)

If the peering between R2 and R5 fails, the new network route will flow from AS100 to AS300 via AS200 instead of flowing directly from AS100 to AS300; as such a simple use of communities can be used to ensure the route is not exported to AS200. You simply need to apply a no-export value to the route as it is advertised on R2 toward R3; this way the route is not advertised to AS200 if a failure occurs. Under normal conditions, AS200 would still see the route from AS300. If you have configured this correctly, as shown in Example 1-25, you have scored 3 points.

EXAMPLE 1-25 Route Advertisement and no-export Configuration on R2

```
R5# sh ip bgp
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
  *>i130.100.200.0/24 120.100.4.1
                   0          100          0 200 10 i
```

```
R2(config)# interface Loopback2
R2(config-if)# ip address 130.100.200.1 255.255.255.0
R2(config-if)# router bgp 10
R2(config-router)# network 130.100.200.0 mask 255.255.255.0
R2(config-router)# neighbor 120.100.3.1 route-map NO-EXPORT out
R2(config-router)# neighbor 120.100.3.1 send-community
R2(config-router)# exit
R2(config)# access-list 5 permit 130.100.200.0
R2(config)# route-map NO-EXPORT permit 10
R2(config-route-map)# match ip address 5
R2(config-route-map)# set community no-export
R2(config-route-map)# route-map NO-EXPORT permit 20

R3# sh ip bgp 130.100.200.1
BGP routing table entry for 130.100.200.0/24, version 4
Paths: (1 available, best #1, table Default-IP-Routing-Table, not advertised to
EBGP peer)
  Advertised to update-groups:
    2
  Local, (Received from a RR-client)
    120.100.2.1 (metric 65) from 120.100.2.1 (130.100.200.1)
      Origin IGP, metric 0, localpref 100, valid, internal, best
      Community: no-export

R5# conf t
Enter configuration commands, one per line. End with CNTL/Z.
R5(config)# int s0/0/1
R5(config-if)# shut
R5(config-if)# ^Z
R5# show ip bgp

R5#
```

- Configure HSRP between R5 and R6 on VLAN300 with R5 active for .1/24. If the network 130.100.200.0/24 is no longer visible to AS300, R6 should dynamically become the HSRP active. Configure R5 to achieve this solution. (4 points)

The clue is in the question; all you need to do is track the specific route with the IP SLA object tracking feature and inform the HSRP process whether the BGP route is withdrawn. You might feel that this isn't strictly a BGP question, but because the IOS section has been removed from the exam, it is possible that topics and features such as this crop up within other sections, so it's best to be aware of as many features as possible.

Because the question doesn't specifically instruct you to configure an exact IP address for your HSRP, you are free to use an unallocated IP address. R5 should be the HSRP active under normal conditions, so this should be configured with the **preempt** command to reinstate control when the route becomes visible once again post withdrawal. Similarly, R6 also requires **preempt** to take control when the priority of R5 decrements. R5 hasn't been configured with a priority in this example because it uses the default value of 100. Example 1-26 shows the configuration and testing steps involved to withdraw the route by shutting down the Frame Relay interface on R5 and toggling the HSRP functionality between R5 and R6. If you have configured this correctly, as shown in Example 1-26, you have scored 4 points.

EXAMPLE 1-26 IP SLA Tracking and HSRP Configuration on R5 and R6

```
R5(config)# track 2 ip route 130.100.200.0 255.255.255.0 reachability
R5(config-track)# interface GigabitEthernet0/1
R5(config-if)# standby 1 ip 150.100.3.1
R5(config-if)# standby 1 preempt

R5(config-if)# standby 1 track 2 decrement 20

R6(config)# interface GigabitEthernet0/1
R6(config-if)# standby 1 ip 150.100.3.1
R6(config-if)# standby 1 priority 90
R6(config-if)# standby 1 preempt

R5# sh standby gigabitEthernet 0/1
GigabitEthernet0/1 - Group 1
  State is Active
    23 state changes, last state change 00:20:11
  Virtual IP address is 150.100.3.1
  Active virtual MAC address is 0000.0c07.ac01
    Local virtual MAC address is 0000.0c07.ac01 (v1 default)
  Hello time 3 sec, hold time 10 sec
  Next hello sent in 0.460 secs
```

```

Preemption enabled
Active router is local
Standby router is 150.100.3.6, priority 90 (expires in 8.472
sec) Priority 100 (default 100)
Track object 2 state Up decrement 20
IP redundancy name is "hsrp-Gi0/1-1" (default)
R5#

R5# conf t
R5(config)# int s0/0/1
R5(config-if)# shut
R5(config-if)#

R5#%BGP-3-NOTIFICATION: sent to neighbor 120.100.2.1 4/0 (hold time expired) 0 bytes
R5#%HSRP-6-STATECHANGE: GigabitEthernet0/1 Grp 1 state Active -> Speak
R5#%HSRP-6-STATECHANGE: GigabitEthernet0/1 Grp 1 state Speak -> Standby
R5# sh standby gigabitEthernet 0/1
GigabitEthernet0/1 - Group 1
  State is Standby
    25 state changes, last state change 00:00:10
  Virtual IP address is 150.100.3.1
  Active virtual MAC address is 0000.0c07.ac01
  Local virtual MAC address is 0000.0c07.ac01 (v1 default)
  Hello time 3 sec, hold time 10 sec
  Next hello sent in 1.880 secs
  Preemption enabled
  Active router is 150.100.3.6, priority 90 (expires in 8.980 sec)
  Standby router is local
  Priority 80 (default 100)
  Track object 2 state Down decrement 20
  IP redundancy name is "hsrp-Gi0/1-1" (default)

```

- Configure two new Loopback interfaces on R1 and R2 of 126.1.1.1/24 and 130.1.1.1/24, respectively, and advertise these into BGP using the **network** command. R3 should be configured to enable only BGP routes originated from R1 up to network 128.0.0.0 and from above network 128.0.0.0 originated from R2. Use only a single ACL on R3 as part of your solution. (3 points)

This is quite an intricate question because you are permitted to use only a single ACL to filter the routes on R3. The method in which you achieve this is to use an ACL that matches networks up to 128.0.0.0 and permits this through one route-map while denying through a separate route-map. The route-maps should be applied on a per-neighbor basis, and both call up the same single ACL. Example 1-27 shows the configuration for the new Loopbacks on R1 and R2 and the filtering on R3. Further testing is detailed in Example 1-28 to substantiate the filtering process on R3. If you have configured this correctly, as shown in Example 1-27, you have scored 3 points.

EXAMPLE 1-27 Route-Map Filtering on R3

```
R1(config)# interface Loopback1
R1(config-if)# ip address 126.1.1.1 255.255.255.0
R1(config-if)# router bgp 10
R1(config-router)# network 126.1.1.0 mask 255.255.255.0

R2(config)# interface Loopback1
R2(config-if)# ip address 130.1.1.1 255.255.255.0
R2(config-if)# router bgp 10
R2(config-router)# network 130.1.1.0 mask 255.255.255.0

R3(config)# access-list 1 permit 0.0.0.0 127.255.255.255
R3(config)# route-map UPTO128 permit 10
R3(config-route-map)# match ip add 1

R3(config)# route-map ABOVE128 permit 10
R3(config-route-map)# match ip add 1
R3(config-route-map)# route-map ABOVE128 permit 20

R3(config)# router bgp 10
R3(config-router)# neighbor 120.100.1.1 route-map UPTO128 in
R3(config-router)# neighbor 120.100.2.1 route-map ABOVE128 in

R3# sh ip bgp
BGP table version is 8, local router ID is 120.100.3.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
```

```

      Network          Next Hop          Metric LocPrf Weight Path
*>i126.1.1.0/24      120.100.1.1          0     100     0  i
*>i130.1.1.0/24      120.100.2.1          0     100     0  i
*>i130.100.200.0/24  120.100.2.1          0     100     0  i

```

```
R3#
```

Further testing of the filtering requires additional interfaces to be configured and advertised on R1 and R2. Example 1-28 shows an interface higher than 128.0.0.0 advertised on R1 and one lower advertised on R2; R3 simply blocks these from entering BGP.

NOTE

This additional testing configuration is not present on the supplied, final configuration.

EXAMPLE 1-28 Route-Map Filtering Verification

```

R1(config)# interface Loopback3
R1(config-if)# ip address 132.1.1.1 255.255.255.0
R1(config-if)# router bgp 10
R1(config-router)# network 132.1.1.0 mask 255.255.255.0
R1(config-router)# ^Z
R1# sh ip bgp neighbors 120.100.3.1 advertised
BGP table version is 7, local router ID is 126.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

```

```

      Network          Next Hop          Metric LocPrf Weight Path
*> 126.1.1.0/24      0.0.0.0           0         32768  i
*> 132.1.1.0/24      0.0.0.0           0         32768  i

```

```
Total number of prefixes 2
```

```

R3# sh ip bgp
BGP table version is 4, local router ID is 120.100.3.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

```

```

      Network          Next Hop          Metric LocPrf Weight Path
*>i126.1.1.0/24      120.100.1.1          0     100     0  i
*>i130.1.1.0/24      120.100.2.1          0     100     0  i
*>i130.100.200.0/24  120.100.2.1          0     100     0  i

```

```

R2# conf t
R2(config)# int Loopback3
R2(config-if)# ip add 100.1.1.1 255.255.255.0
R2(config-if)# router bgp 10
R2(config-router)# network 100.1.1.0 mask 255.255.255.0
R2(config-router)# ^Z
R2# sh ip bgp neighbor 120.100.3.1 advertised
BGP table version is 5, local router ID is 130.100.200.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network                Next Hop           Metric LocPrf Weight Path
*> 100.1.1.0/24           0.0.0.0             0         32768 i
*> 130.1.1.0/24           0.0.0.0             0         32768 i
*> 130.100.200.0/24      0.0.0.0             0         32768 i

Total number of prefixes 3

R3# sh ip bgp
BGP table version is 4, local router ID is 120.100.3.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network                Next Hop           Metric LocPrf Weight Path
*>i126.1.1.0/24          120.100.1.1         0      100      0 i
*>i130.1.1.0/24          120.100.2.1         0      100      0 i
*>i130.100.200.0/24     120.100.2.1         0      100      0 i

```

Section 4: IPv6 (14 Points)

The prerequisite to the questions is configuration of the IPv6 addresses and Frame Relay. You should test your IPv6 connectivity to ensure you are ready to progress to the routing questions. You will of course need Frame Relay maps to achieve connectivity. Unlike IPv4, though, you will need two maps, one to reach the IPv6 remote address over the PVC and one to map to the remote Link Local addresses. Example 1-29 shows the initial testing over Frame Relay and

required IPv6 configuration to progress to the routing questions. Consider using the **show ipv6 interfaces brief** command for a quick check of your interface configuration.

EXAMPLE 1-29 IPv6 Testing and Initial Configuration

```
R1# debug frame-relay packet
Frame Relay packet debugging is on
R1# ping ipv6 2007:C15:C0:11::3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2007:C15:C0:11::3, timeout is 2 seconds:

Serial0/0/0:Encaps failed--no map entry link 79(IPV6)

R1# conf t
R1(config)# int s0/0/0
R1(config-if)# frame-relay map ipv6 2007:C15:C0:11::3 103 broadcast
R1(config-if)# ^Z
R1#

R3# conf t
R3(config)# int s0/0/0
R3(config-if)# frame-relay map ipv6 2007:C15:C0:11::1 301 broadcast
R3(config-if)# ^Z

R1# ping ipv6 2007:C15:C0:11::3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2007:C15:C0:11::3, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 4/4/4 ms
R1#

R3# sh ipv6 int s0/0/0 | include link-local
IPv6 is enabled, link-local address is FE80::214:6AFF:FEFC:7390
No Virtual link-local address(es):
R3#
```

```
R1# sh ipv6 interface s0/0/0 | include link-local
IPv6 is enabled, link-local address is FE80::213:C3FF:FE7B:E3C0
No Virtual link-local address(es):
R1#
```

```
R2# sh ipv6 interface s0/0 | include link-local
IPv6 is enabled, link-local address is FE80::213:7FFF:FE84:BEE0
No Virtual link-local address(es):
```

```
R2# sh ipv6 interface s0/1 | include link-local
IPv6 is enabled, link-local address is FE80::213:7FFF:FE84:BEE0
No Virtual link-local address(es):
R2#
```

```
R5# sh ipv6 interface s0/0/1 | include link-local
IPv6 is enabled, link-local address is FE80::214:6AFF:FEFC:F130
No Virtual link-local address(es):
R5#
```

```
R1(config)# ipv6 unicast-routing
R1(config)# interface gigabitEthernet 0/1
R1(config-if)# ipv6 address 2007:C15:C0:10::1/64
R1(config-if)# interface Serial0/0/0
R1(config-if)# ipv6 address 2007:C15:C0:11::1/64
R1(config-if)# frame-relay map ipv6 2007:C15:C0:11::3 103 broadcast
R1(config-if)# frame-relay map ipv6 2007:C15:C0:11::2 103 broadcast
R1(config-if)# frame-relay map ipv6 FE80::213:7FFF:FE84:BEE0 103 broadcast
R1(config-if)# frame-relay map ipv6 FE80::214:6AFF:FEFC:7390 103 broadcast
```

```
R2(config)# ipv6 unicast-routing
R2(config)# interface fastEthernet 0/1
R2(config-if)# ipv6 address 2007:C15:C0:12::2/64
R2(config-if)# interface serial 0/0
R2(config-if)# ipv6 address 2007:C15:C0:11::2/64
R2(config-if)# frame-relay map ipv6 2007:C15:C0:11::1 203 broadcast
R2(config-if)# frame-relay map ipv6 2007:C15:C0:11::3 203 broadcast
R2(config-if)# frame-relay map ipv6 2007:C15:C0:10::12 203 broadcast
R2(config-if)# frame-relay map ipv6 FE80::213:C3FF:FE7B:E3C1 203 broadcast
```

```
R2(config-if)# frame-relay map ipv6 FE80::214:6AFF:FEFC:7390 203 broadcast
R2(config-if)# interface serial 0/1
R2(config-if)# ipv6 address 2007:C15:C0:14::2/64
R2(config-if)# frame-relay map ipv6 FE80::214:6AFF:FEFC:F130 215 broadcast
R2(config-if)# frame-relay map ipv6 2007:C15:C0:14::5 215 broadcast
```

```
R3(config)# ipv6 unicast-routing
R3(config)# interface gigabitEthernet 0/0
R3(config-if)# ipv6 address 2007:C15:C0:15::3/64
R3(config-if)# interface serial 0/0/0
R3(config-if)# ipv6 address 2007:C15:C0:11::3/64
R3(config-if)# frame-relay map ipv6 2007:C15:C0:11::1 301 broadcast
R3(config-if)# frame-relay map ipv6 2007:C15:C0:11::2 302 broadcast
```

```
R4(config)# ipv6 unicast-routing
R4(config)# interface gigabitEthernet 0/0
R4(config-if)# ipv6 address 2007:C15:C0:15::4/64
```

```
R5(config)# ipv6 unicast-routing
R5(config)# interface gigabitEthernet 0/1
R5(config)# ipv6 address 2007:C15:C0:16::5/64
R5(config-if)# interface Serial0/0/1
R5(config-if)# ipv6 address 2007:C15:C0:14::5/64
R5(config-if)# frame-relay map ipv6 2007:C15:C0:14::2 512 broadcast
R5(config-if)# frame-relay map ipv6 FE80::213:7FFF:FE84:BEE0 512 broadcast
```

```
R6(config)# ipv6 unicast-routing
R6(config)# interface gigabitEthernet 0/1
R6(config-if)# ipv6 address 2007:C15:C0:16::6/64
```

Section 4.1: RIPng

- Configure RIPng ensuring your IPv6 routes are visible throughout your RIPng domain. Do not disable split-horizon. (3 points)

R3 by default has split horizon enabled on the Frame Relay interface; the hub receives both R1 and R2 Ethernet associated IPv6 routes but because of split-horizon will not advertise these back out onto the same interface. As you are not permitted to disable split-horizon, you will need to create a tunnel between R1 and R2. Example 1-30 shows the initial RIPng configuration and routing tables of R1 and R2 without each other's Ethernet IPv6 routes present and the required tunnel configuration. If you have configured this correctly, as shown in Example 1-30, you have scored 3 points.

EXAMPLE 1-30 RIPng Configuration and Testing

```
R1(config)# interface gigabitEthernet 0/1
R1(config-if)# ipv6 rip CCIE enable
R1(config-if)# interface Serial0/0/0
R1(config-if)# ipv6 rip CCIE enable
```

```
R2(config)# interface fastEthernet 0/1
R2(config-if)# ipv6 rip CCIE enable
R2(config-if)# interface serial 0/0
R2(config-if)# ipv6 rip CCIE enable
R2(config-if)# interface serial 0/1
R2(config-if)# ipv6 rip CCIE enable
```

```
R3(config)# interface gigabitEthernet 0/0
R3(config-if)# ipv6 rip CCIE enable
R3(config-if)# interface serial 0/0/0
R3(config-if)# ipv6 rip CCIE enable
```

```
R4(config)# interface gigabitEthernet 0/0
R4(config-if)# ipv6 rip CCIE enable
```

```
R5(config)# interface Serial0/0/1
R5(config-if)# ipv6 rip CCIE enable
```

```
R1# show ipv6 route rip
```

```
IPv6 Routing Table - 10 entries
```

```
Codes: C - Connected, L - Local, S - Static, R - RIP, B -
```

```
        BGP U - Per-user Static route
```

```
        I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
```

```
        O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
```

```
ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
D - EIGRP, EX - EIGRP external
R 2007:C15:C0:15::/64 [120/2]
   via FE80::214:6AFF:FEFC:7390, Serial0/0/0
```

R2# show ipv6 route rip

```
IPv6 Routing Table - 10 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B -
       BGP U - Per-user Static route
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
       D - EIGRP, EX - EIGRP external
R 2007:C15:C0:15::/64 [120/2]
   via FE80::214:6AFF:FEFC:7390, Serial0/0
```

R3# show ipv6 route rip

```
IPv6 Routing Table - 10 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B -
       BGP U - Per-user Static route
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
       D - EIGRP, EX - EIGRP external
R 2007:C15:C0:10::/64 [120/2]
   via FE80::213:C3FF:FE7B:E3C0, Serial0/0/0
R 2007:C15:C0:12::/64 [120/2]
   via FE80::213:7FFF:FE84:BEE0, Serial0/0/0
```

R1(config)# interface Tunnel1

```
R1(config-if)# ipv6 address 2007:C15:C0:13::1/64
R1(config-if)# ipv6 rip CCIE enable
R1(config-if)# tunnel source Serial0/0/0
R1(config-if)# tunnel destination 120.100.123.2
R1(config-if)# tunnel mode ipv6ip
```

R2(config)# interface Tunnel1

```
R2(config-if)# ipv6 address 2007:C15:C0:13::2/64
```

```
R2(config-if)# ipv6 rip CCIE enable
R2(config-if)# tunnel source Serial0/0
R2(config-if)# tunnel destination 120.100.123.1
R2(config-if)# tunnel mode ipv6ip
```

```
R1# show ipv6 route rip
```

```
IPv6 Routing Table - 11 entries
```

```
Codes: C - Connected, L - Local, S - Static, R - RIP, B -
```

```
       BGP U - Per-user Static route
```

```
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
```

```
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
```

```
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
```

```
       D - EIGRP, EX - EIGRP external
```

```
R   2007:C15:C0:12::/64 [120/2]
```

```
R   2007:C15:C0:15::/64 [120/2]
```

```
    via FE80::214:6AFF:FEFC:7390, Serial0/0/0
```

```
R2# show ipv6 route rip
```

```
IPv6 Routing Table - 13 entries
```

```
Codes: C - Connected, L - Local, S - Static, R - RIP, B -
```

```
       BGP U - Per-user Static route
```

```
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
```

```
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
```

```
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
```

```
       D - EIGRP, EX - EIGRP external
```

```
R   2007:C15:C0:10::/64 [120/2]
```

```
    via FE80::7864:7B01, Tunnell
```

```
R   2007:C15:C0:15::/64 [120/2]
```

```
    via FE80::214:6AFF:FEFC:7390, Serial0/0
```

Section 4.2: OSPFv3

- Configure OSPFv3 with a process ID of 1 with all OSPF interfaces assigned to area 0. (2 points).

This is a clear-cut OSPFv3 configuration. If you have configured this correctly, as shown in Example 1-31, you have scored 2 points.

EXAMPLE 1-31 R5 and R6 OSPFv3 Configuration

```
R5(config)# interface gigabitEthernet 0/1
R5(config-if)# ipv6 ospf 1 area 0
```

```
R6(config)# interface gigabitEthernet 0/1
R6(config-if)# ipv6 ospf 1 area 0
```

```
R5# show ipv6 ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Interface ID	Interface
120.100.6.1	1	FULL/DR	00:00:30	3	GigabitEthernet0/1

```
R6# show ipv6 ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Interface ID	Interface
120.100.5.1	1	FULL/BDR	00:00:39	3	GigabitEthernet0/1

- The IPv6 network is deemed to be stable; as such, reduce the number of LSAs flooded within the OSPF domain. (2 points)

To suppress the unnecessary flooding of link-state advertisements in stable topologies, the **ipv6 ospf flood-reduction** command is required under interface configuration mode. If you have configured this correctly, as shown in Example 1-32, you have scored 2 points.

EXAMPLE 1-32 R5 and R6 Flood-Reduction Configuration

```
R5(config)# interface gigabitEthernet 0/1
R5(config-if)# ipv6 ospf flood-reduction
```

```
R6(config)# interface gigabitEthernet 0/1
R6(config-if)# ipv6 ospf flood-reduction
```

Section 4.3: Redistribution

- Redistribute RIPng routes into the OSPFv3 demand (one way); RIP routes should have a fixed cost of 5000 associated to them within the OSPF network. (1 point)

As per vanilla OSPF, the default behavior for OSPFv3 is for redistributed routes to be advertised with a fixed cost as type 2 external routes, so a simple redistribution configuration with a default-metric of 5000 on R5 is required. Example 1-33 shows the required configuration and routing table on R6 for the redistributed RIPng routes. Pay attention to ensure you have full route visibility because the Frame Relay network on R5 (2007:C15:C0:14::) will not be present within the OSPFv3 domain unless R5 specifically redistributes its own connected interfaces. If you have configured this correctly, as shown in Example 1-33, you have scored 1 point.

EXAMPLE 1-33 R5 OSPFv3 Redistribution Configuration

```
R5(config)# ipv6 router ospf 1
R5(config-router)# redistribute rip CCIE metric 5000

R6# sh ipv6 route ospf
IPv6 Routing Table - 10 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B -
       BGP U - Per-user Static route
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
       D - EIGRP, EX - EIGRP external
OE2 2007:C15:C0:10::/64 [110/5000]
    via FE80::214:6AFF:FEFC:F131, GigabitEthernet0/1
OE2 2007:C15:C0:11::/64 [110/5000]
    via FE80::214:6AFF:FEFC:F131, GigabitEthernet0/1
OE2 2007:C15:C0:12::/64 [110/5000]
    via FE80::214:6AFF:FEFC:F131, GigabitEthernet0/1
OE2 2007:C15:C0:13::/64 [110/5000]
    via FE80::214:6AFF:FEFC:F131, GigabitEthernet0/1
OE2 2007:C15:C0:15::/64 [110/5000]
    via FE80::214:6AFF:FEFC:F131, GigabitEthernet0/1

R5(config)# ipv6 router ospf 1
R5(config-rtr)# redistribute rip CCIE metric 5000 include-connected
```

```

R6# show ipv6 route 2007:C15:C0:14::
IPv6 Routing Table - 10 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B -
       BGP U - Per-user Static route
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
       D - EIGRP, EX - EIGRP external
OE2 2007:C15:C0:14::/64 [110/5000]
    via FE80::214:6AFF:FEFC:F131, GigabitEthernet0/1

```

- Ensure the OSPF3 network is reachable from the RIP network by a single route of 2007::/16, which should be seen within the RIP domain. Configure R5 only to achieve this. The OSPF domain should continue to receive specific RIPng subnets. (2 points)

As you are not mutually redistributing protocols, you are required to configure an IPv6 summary route into the RIPng domain on R5 to provide full connectivity from the RIPng domain into OSPFv3. If you have configured this correctly, as shown in Example 1-34, you have scored 2 points.

EXAMPLE 1-34 R5 RIPng Summary Configuration and Connectivity Testing

```

R5(config-if)# int s0/0/1
R5(config-if)# ipv6 rip CCIE summary-address 2007::/16

R1# show ipv6 route rip
IPv6 Routing Table - 13 entries
R 2007::/16 [120/3]
  via FE80::213:7FFF:FE84:BEE0, Tunnell
R 2007:C15:C0:12::/64 [120/2]
  via FE80::213:7FFF:FE84:BEE0, Tunnell
R 2007:C15:C0:14::/64 [120/2]
  via FE80::213:7FFF:FE84:BEE0, Tunnell
R 2007:C15:C0:15::/64 [120/2]
  via FE80::214:6AFF:FEFC:7390, Serial10/0/0

R1#
R1# ping ipv6 2007:C15:C0:16::5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2007:C15:C0:16::5, timeout is 2 seconds:

```

```

!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 12/12/16 ms
R1# ping ipv6 2007:C15:C0:16::6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2007:C15:C0:16::6, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 12/15/16 ms

```

- Ensure if the serial link fails between the OSPF and RIPng domain that routing is still possible between R5 and R4 over VLAN45. Do not enable RIP on the VLAN45 interfaces of R4 and R5—configure R4 and R5 to achieve this, and this should be considered as an alternative path only if a failure occurs. (3 points)

R4 and R5 both belong to the RIPng domain. If you can't enable RIPng on the VLAN45 interfaces, all you can do is create a tunnel between the devices. You might have considered enabling OSPFv3 between routers, but you have not been given sufficient information to perform this, and it would then create additional problems in terms of redistribution points. Example 1-35 shows the required configuration to tunnel IPv6 through IPv4 on R4 and R5. You should notice that certain routes will have a lower hop count through the tunnel as opposed to through the physical RIPng network. The question states that the newly configured link should be used only if a failure occurs. As such, you need to penalize the tunnel by use of an **offset-list** applied directly to the tunnel interface of R4 and R5. R5 will still receive the summary /16 route configured earlier via the tunnel regardless of how high you set the hop count. The following question addresses this condition. If you have configured this correctly, as shown in Example 1-35, you have scored 3 points.

EXAMPLE 1-35 R4 and R5 Tunnel Configuration and Verification

```

R4(config)# interface Tunnel0
R4(config-if)# ipv6 address 2007:C15:C0:17::4/64
R4(config-if)# ipv6 rip CCIE enable
R4(config-if)# tunnel source GigabitEthernet0/1.45
R4(config-if)# tunnel destination 120.100.45.5
R4(config-if)# tunnel mode ipv6ip

R5(config)# interface Tunnel0
R5(config-if)# ipv6 address 2007:C15:C0:17::5/64
R5(config-if)# ipv6 rip CCIE enable
R5(config-if)# tunnel source GigabitEthernet0/0
R5(config-if)# tunnel destination 120.100.45.4
R5(config-if)# tunnel mode ipv6ip

```

R4# **show ipv6 route rip**

IPv6 Routing Table - 12 entries

Codes: C - Connected, L - Local, S - Static, R - RIP, B -

BGP U - Per-user Static route

I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary

O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2

ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2

D - EIGRP, EX - EIGRP external

```
R 2007::/16 [120/4]
   via FE80::214:6AFF:FEFC:7390, GigabitEthernet0/0
R 2007:C15:C0:10::/64 [120/3]
   via FE80::214:6AFF:FEFC:7390, GigabitEthernet0/0
R 2007:C15:C0:11::/64 [120/2]
   via FE80::214:6AFF:FEFC:7390, GigabitEthernet0/0
R 2007:C15:C0:12::/64 [120/3]
   via FE80::214:6AFF:FEFC:7390, GigabitEthernet0/0
   via FE80::7864:2D05, Tunnel0
R 2007:C15:C0:13::/64 [120/3]
   via FE80::214:6AFF:FEFC:7390, GigabitEthernet0/0
   via FE80::7864:2D05, Tunnel0
R 2007:C15:C0:14::/64 [120/2]
   via FE80::7864:2D05, Tunnel0
```

R5# **show ipv6 route rip**

IPv6 Routing Table - 14 entries

Codes: C - Connected, L - Local, S - Static, R - RIP, B -

BGP U - Per-user Static route

I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary

O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2

ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2

D - EIGRP, EX - EIGRP external

```
R 2007::/16 [120/5]
   via FE80::7864:2D04, Tunnel0
R 2007:C15:C0:10::/64 [120/3]
   via FE80::213:7FFF:FE84:BEE0, Serial0/0/1
R 2007:C15:C0:11::/64 [120/2]
   via FE80::213:7FFF:FE84:BEE0, Serial0/0/1
R 2007:C15:C0:12::/64 [120/2]
   via FE80::213:7FFF:FE84:BEE0, Serial0/0/1
```

```
R 2007:C15:C0:13::/64 [120/2]
    via FE80::213:7FFF:FE84:BEE0, Serial0/0/1
R 2007:C15:C0:15::/64 [120/2]
    via FE80::7864:2D04, Tunnel0

R5(config)# interface Tunnel0
R5(config-if)# ipv6 rip CCIE metric-offset 4
R5(config-if)# do show ipv6 route rip
IPv6 Routing Table - 14 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B -
        BGP U - Per-user Static route
        I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
        O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
        ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
        D - EIGRP, EX - EIGRP external
R 2007::/16 [120/8]
    via FE80::7864:2D04, Tunnel0
R 2007:C15:C0:10::/64 [120/3]
    via FE80::213:7FFF:FE84:BEE0, Serial0/0/1
R 2007:C15:C0:11::/64 [120/2]
    via FE80::213:7FFF:FE84:BEE0, Serial0/0/1
R 2007:C15:C0:12::/64 [120/2]
    via FE80::213:7FFF:FE84:BEE0, Serial0/0/1
R 2007:C15:C0:13::/64 [120/2]
    via FE80::213:7FFF:FE84:BEE0, Serial0/0/1
R 2007:C15:C0:15::/64 [120/3]
    via FE80::213:7FFF:FE84:BEE0, Serial0/0/1

R4(config)# interface Tunnel0
R4(config-if)# ipv6 rip CCIE metric-offset 4
R4(config-if)# do show ipv6 route rip
IPv6 Routing Table - 12 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B -
        BGP U - Per-user Static route
        I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
        O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
        ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
        D - EIGRP, EX - EIGRP external
R 2007::/16 [120/4]
    via FE80::214:6AFF:FEFC:7390, GigabitEthernet0/0
```

```

R 2007:C15:C0:10::/64 [120/3]
   via FE80::214:6AFF:FEFC:7390, GigabitEthernet0/0
R 2007:C15:C0:11::/64 [120/2]
   via FE80::214:6AFF:FEFC:7390, GigabitEthernet0/0
R 2007:C15:C0:12::/64 [120/3]
   via FE80::214:6AFF:FEFC:7390, GigabitEthernet0/0
R 2007:C15:C0:13::/64 [120/3]
   via FE80::214:6AFF:FEFC:7390, GigabitEthernet0/0
R 2007:C15:C0:14::/64 [120/3]
   via FE80::214:6AFF:FEFC:7390, GigabitEthernet0/0

```

- Ensure that the summary route configured previously is not seen back on the routing table of R5. Configure only R5 to achieve this. (1 point)

As briefly discussed in the previous question, the summary route will return to R5 through the newly created tunnel interface. This is expected behavior because of the method in which it was originally advertised. A simple **prefix-list** is required on R5 to deny the summary and permit all other routes entering the tunnel interface. If you have configured this correctly, as shown in Example 1-36, you have scored 3 points.

EXAMPLE 1-36 R5 Distribute-list Configuration and Verification

```

R5# show ipv6 route rip
IPv6 Routing Table - 14 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B -
       BGP U - Per-user Static route
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
       D - EIGRP, EX - EIGRP external
R 2007::/16 [120/8]
   via FE80::7864:2D04, Tunnel0
R 2007:C15:C0:10::/64 [120/3]
   via FE80::213:7FFF:FE84:BEE0, Serial0/0/1
R 2007:C15:C0:11::/64 [120/2]
   via FE80::213:7FFF:FE84:BEE0, Serial0/0/1
R 2007:C15:C0:12::/64 [120/2]
   via FE80::213:7FFF:FE84:BEE0, Serial0/0/1
R 2007:C15:C0:13::/64 [120/2]
   via FE80::213:7FFF:FE84:BEE0, Serial0/0/1
R 2007:C15:C0:15::/64 [120/3]

```

```
via FE80::213:7FFF:FE84:BEE0, Serial0/0/1

R5(config)# ipv6 prefix-list BLOCK-SUMMARY seq 10 deny 2007::/16
R5(config)# ipv6 prefix-list BLOCK-SUMMARY seq 15 permit ::/0 le 128
R5(config)# ipv6 router rip CCIE
R5(config-router)# distribute-list prefix-list BLOCK-SUMMARY in Tunnel0
R5(config-router)# do show ipv6 route rip
IPv6 Routing Table - 13 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B -
       BGP U - Per-user Static route
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
       D - EIGRP, EX - EIGRP external
R 2007:C15:C0:10::/64 [120/3]
   via FE80::213:7FFF:FE84:BEE0, Serial0/0/1
R 2007:C15:C0:11::/64 [120/2]
   via FE80::213:7FFF:FE84:BEE0, Serial0/0/1
R 2007:C15:C0:12::/64 [120/2]
   via FE80::213:7FFF:FE84:BEE0, Serial0/0/1
R 2007:C15:C0:13::/64 [120/2]
   via FE80::213:7FFF:FE84:BEE0, Serial0/0/1
R 2007:C15:C0:15::/64 [120/3]
   via FE80::213:7FFF:FE84:BEE0, Serial0/0/1
```

Section 5: QoS (8 Points)

- You are required to configure QoS on Switch1 according to the Cisco QoS baseline model. Create a Modular QoS configuration for all user ports (Fast Ethernet 1-24) that facilitates the following requirements (3 points):
 - 1) All ports should trust the DSCP values received from their connecting devices.
 - 2) Packets received from the user ports with DSCP values of 48, 46, 34, 32, 24, 28, 16 and 10 should be remarked to DSCP 8 (PHB CS1) *if* traffic flowing occurs above 5 Mbps on a per port basis. This traffic could be a combination of any of the preceding DSCP values with any source/destination combination. Ensure a minimum burst value is configured above the 5 Mbps.

It is acknowledged within the industry that a user port rarely generates more than 5 Mbps of traffic on a standard FastEthernet connection. If traffic rates increase above this threshold, it could be indicative of a DOS or Worm attack. A

method of mitigating an attack is to create a Scavenger-Class that simply remarks traffic DSCP values when the threshold has been exceeded. This will not block traffic but will ensure that mission-critical traffic remains unaffected from an attack by trusting the DSCP value for known traffic and re-marking unknown application traffic down to CS1.

To answer the question, you are required to create a Modular QoS policy that trusts the incoming DSCP value received from the host within the policy rather than by configuring the trust value on a per-interface basis and by policing traffic at a rate of 5 Mbps. When the minimum burst rate is exceeded, the DSCP values will be remapped according to the **policed-dscp** map to Scavenger-Class CS1 (DSCP8). You should note that all DSCP baseline values are being remapped with the exception of DSCP26, which is generally reserved for mission-critical data. This approach enables traffic associated with this value to remain unchanged even when traffic rates exceed 5 Mbps; this approach also assumes that the virus does not itself re-mark traffic to this value to increase its chances of causing damage. The exclusion of DSCP26 though is not relevant to the configuration and methodology you use to answer the question. The question requires you to configure a standard IP ACL that permits any traffic. For traffic matching this classification, the DSCP value in the incoming packet is trusted. If the matched traffic exceeds an average traffic rate of 5 Mbps and a normal burst size of 8000 bytes, its DSCP is marked down according to the policed DSCP map values and transmitted. If you have configured this correctly, as shown in Example 1-37, you have scored 3 points.

EXAMPLE 1-37 R5 Distribute-list Configuration and Verification

```
SW1(config)# mls qos
SW1(config)# mls qos map policed-dscp 48 46 34 32 24 28 16 10 to 8
SW1(config)# access-list 1 permit any
SW1(config)# class-map POLICE
SW1(config-cmap)# match access-group 1
SW1(config-cmap)# exit
SW1(config)# policy-map RE-MARK
SW1(config-pmap)# class POLICE
SW1(config-pmap-c)# trust dscp
SW1(config-pmap-c)# police 5000000 8000 exceed-action policed-dscp-transmit
SW1(config-pmap-c)# exit
SW1(config-pmap)# exit
SW1(config)# interface range fastEthernet 0/1-24
SW1(config-if-range)# service-policy input RE-MARK

SW1# show policy-map RE-MARK
```

```

Policy Map RE-MARK
Class POLICE
  police 5000000 8000 exceed-action policed-dscp-transmit
  trust dscp
    
```

- Switch1 will be connected to a new trusted domain in the future using interface gigabit 0/1. A DSCP value received locally on sw1 of AF43 should be mapped to AF42 when destined for the new domain. (2 points)

This requires a DSCP mutation map to convert DSCP values between environments. If you didn't realize that AF43 is DSCP38 and AF42 is DSCP36, you would struggle to answer this question, but a search of your documentation CD should have assisted you. For the mutation map to function correctly, you need to explicitly trust DSCP values received on the interface on which you are configuring the map. If you have configured this correctly, as shown in Example 1-38, you have scored 2 points.

EXAMPLE 1-38 Switch1 DSCP-mutation Map Configuration

```

SW1(config)# mls qos map dscp-mutation AF43-TO-AF42 38 to 36
SW1(config)# interface Gig0/1
SW1(config-if)# mls qos trust dscp
SW1(config-if)# mls qos dscp-mutation AF43-TO-AF42
    
```

- Configure Cisco Modular QoS as follows on R1 for the following traffic types based on their associated Per Hop Behavior into classes. Incorporate these into an overall policy that should be applied to the T1 interface S0/0/0. Assume a PVC of line rate on the Frame Relay network, and allow each class the effective bandwidth as detailed (2 points).

Class	PHB	Assigned Speed
Routing	CS6	46 Kbps
VoIP	EF	247 Kbps
Interactive Video	AF41	247 Kbps
Mission Critical Data	AF31	247 Kbps
Call-Signaling	CS3	46 Kbps
Transactional Data	AF21	216 Kbps
Network-mgmt	CS2	46 Kbps
Bulk Data	Af11	46 Kbps

Scavenger	CS1	15 Kbps
Default	0	386 Kbps

Two points are available here, so you know it's either going to be complex or involve a great deal of configuration. This one is a bit of both, so there is a risk of configuration errors for those points to slip away. There is also some math involved because the **policy-map** requires a percentage value of bandwidth as opposed to actual speed, as you are using a T1 interface you know that the maximum available bandwidth is 1544 Kbs and a line rate PVC is assumed, so the values required are as follows:

1% = 15 Kbps, 3% = 46 Kbps, 14% = 216 Kbps, 16% = 247 Kbps, 25% = 386 Kbps.

A **class-map** to match all values for the provided classes is required that is then associated with the **policy-map**. The overall policy is then applied to the outgoing interface Serial0/0/0, and a nice little gotcha is that you must configure the interface with the command **max-reserved-bandwidth 100**; otherwise, the full bandwidth is not made available for the policy. Usually you would assign voice traffic into a real-time queue (LLQ), but the question doesn't dictate this, so effectively all traffic types are being assigned with different proportions of CBWFQ. If you have configured this correctly, as shown in Example 1-39, you have scored 2 points.

EXAMPLE 1-39 Switch1 Modular QoS Configuration

```
R1# sh run class-map
!
class-map match-all VOIP
  match ip dscp ef
class-map match-all BULK-DATA
  match ip dscp af11
class-map match-all NET-MAN
  match ip dscp cs2
class-map match-all VIDEO
  match ip dscp af41
class-map match-all ROUTING
  match ip dscp cs6
class-map match-all SCAVENGER
  match ip dscp cs1
class-map match-all TRANS-DATA
  match ip dscp af21
class-map match-all MISSION-CRIT
  match ip dscp af31
```

```
class-map match-all CALL-SIG
  match ip dscp cs3
!
end
```

```
R1# sh run policy-map
!
policy-map QOS
  class VOIP
    bandwidth percent 16
  class VIDEO
    bandwidth percent 16
  class BULK-DATA
    bandwidth percent 3
    random-detect
  class TRANS-DATA
    bandwidth percent 14
  class NET-MAN
    bandwidth percent 3
  class ROUTING
    bandwidth percent 3
  class SCAVENGER
    bandwidth percent 1
  class MISSION-CRIT
    bandwidth percent 16
    random-detect
  class CALL-SIG
    bandwidth percent 3
  class class-default
    bandwidth percent 25
!
end
```

```
R1# sh run int s0/0/0 | begin max-reserved-bandwidth 100
max-reserved-bandwidth 100
service-policy output QOS
end
```

```
R1# show policy-map QOS
Policy Map QOS
```

```

Class VOIP
  Bandwidth 16 (%) Max Threshold 64 (packets)
Class VIDEO
  Bandwidth 16 (%) Max Threshold 64 (packets)
Class BULK-DATA
  Bandwidth 3 (%)
    exponential weight 9
    class      min-threshold    max-threshold    mark-probability
    -----
    0          -                -                1/10
    1          -                -                1/10
    2          -                -                1/10
    3          -                -                1/10
    4          -                -                1/10
    5          -                -                1/10
    6          -                -                1/10
    7          -                -                1/10
    rsvp      -                -                1/10

Class TRANS-DATA
  Bandwidth 14 (%) Max Threshold 64 (packets)
Class NET-MAN
  Bandwidth 3 (%) Max Threshold 64 (packets)
Class ROUTING
  Bandwidth 3 (%) Max Threshold 64 (packets)
Class SCAVENGER
  Bandwidth 1 (%) Max Threshold 64 (packets)
Class MISSION-CRIT
  Bandwidth 16 (%)
    exponential weight 9
    class      min-threshold    max-threshold    mark-probability
    -----
    0          -                -                1/10
    1          -                -                1/10
    2          -                -                1/10
    3          -                -                1/10
    4          -                -                1/10
    5          -                -                1/10
  
```

```

        6          -          -          1/10
        7          -          -          1/10
        rsvp       -          -          1/10

Class CALL-SIG
  Bandwidth 3 (%) Max Threshold 64 (packets)
Class class-default
  Bandwidth 25 (%)
    exponential weight 9
    class      min-threshold    max-threshold    mark-probability
-----
        0          -          -          1/10
        1          -          -          1/10
        2          -          -          1/10
        3          -          -          1/10
        4          -          -          1/10
        5          -          -          1/10
        6          -          -          1/10
        7          -          -          1/10
        rsvp       -          -          1/10

```

- Configure R2 so that traffic can be monitored on the Frame Relay network with a view to a dynamic policy being generated in the future that trusts the DSCP value of traffic identified on this media. (1 point)

This is a simple question that requires the command **auto discovery qos trust** be configured under the Frame Relay interface of R2. This command uses NBAR to inspect the application traffic that flows through the router with a view of generating a QoS policy based on the traffic flow profile. The keyword **trust** in the command ensures that the DSCP value of the traffic monitored on the network is trusted. If you have configured this correctly, you have scored 1 point.

Section 6: Security (6 Points)

- Configure R3 to identify and discard the following custom virus. The virus is characterized by the ASCII characters `Hastings_Beer` within the payload and utilizes UDP ports 11664 to 11666. The ID of the virus begins on the third character of the payload. The virus originated on VLAN 34. (4 points)

This fictitious virus requires the use of NBAR with PDLM to inspect a packet payload to identify the virus based on the information supplied within the question. As the virus is located within the third ASCII character, you need to inform the custom NBAR list to ignore the first two characters, which ensures that it will begin to check the third packet. If you have configured this correctly, as shown in Example 1-40, you have scored 3 points. You can use the **show policy-map** command to verify your configuration.

EXAMPLE 1-40 R3 NBAR Configuration

```
R3(config)# ip nbar custom Hastings_Beer 2 ascii Hastings_Beer udp range 11664 11666
R3(config)# class-map match-all VIRUS
R3(config-cmap)# match protocol Hastings_Beer
R3(config-cmap)# policy-map BLOCK-VIRUS
R3(config-pmap)# class VIRUS
R3(config-pmap-c)# drop
R3(config-pmap-c)# interface gigabit0/0
R3(config-if)# Service-policy input BLOCK-VIRUS
```

- There is an infected host on VLAN 200 of 150.100.2.100. Ensure that only within BGP AS10, traffic destined for this host is directed to null0 of each local router. You may not use any ACLs to block traffic to this host specifically but may use a static route pointing to null 0 for traffic destined to 192.0.2.0 /24 on routers within AS10. R2 may have an additional static route pointing to null0. Use a BGP feature on R2 to ensure traffic to this source is blocked. Prevent unnecessary replies when traffic is passed to the null0 interface for users residing on VLAN100. (4 points)

This question is representative of black-hole routing. This is an effective method of discarding packets being sent to a known destination. This approach to discarding traffic is efficient because it enables the edge routers to route traffic rather than use ACLs, and it can be deployed dynamically by making use of the next-hop field within BGP updates. You are permitted to create a static route on Routers R1, R2, and R3 in AS10 for network 192.0.2.0/24 to null0 and one additional route on R2. This route would need to be directing traffic to the infected host to null0, to update Routers R1 and R3. R2 simply advertises the host route for the infected host to AS10 and sets the next-hop for this to 192.0.2.1. Routers R1 and R3 then direct traffic to null0 when traffic is destined to the infected host. To ensure the solution is only used in AS10, you need to set the community to **no-export** for the specific static route and tag the route with a value of 10 to identify it. You must therefore send the community values to neighbor R3 on R2, but this should have completed previously for an earlier BGP question. Use of the **no icmp unreachable** command on R1's GigabitEthernet interface prevents unnecessary replies when traffic is passed to the Null0 interface. If you have configured this correctly, as shown in Example 1-41, you have scored 3 points.

EXAMPLE 1-41 BGP Black Hole Routing Configuration and Verification

```
R2(config)# ip route 192.0.2.1 255.255.255.255 null0
R2(config)# ip route 150.100.2.100 255.255.255.255 Null0 Tag 10
R2(config)# router bgp 10
R2(config-router)# redistribute static route-map BLACKHOLE
R2(config-router)# route-map BLACKHOLE permit 10
R2(config-route-map)# match tag 10
R2(config-route-map)# set ip next-hop 192.0.2.1
R2(config-route-map)# set community no-export
R2(config-route-map)# exit
R2(config)# ip route 192.0.2.1 255.255.255.255 null0
R2(config)# do show ip bgp neigh 120.100.3.1 advertised
BGP table version is 6, local router ID is 130.100.200.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
Network          Next Hop          Metric LocPrf Weight Path
*> 130.1.1.0/24   0.0.0.0           0         32768 i
*> 130.100.200.0/24 0.0.0.0           0         32768 i
*> 150.100.2.100/32 192.0.2.1         0         32768 i
Total number of prefixes 3

R2# show ip route 150.100.2.100
Routing entry for 150.100.2.100/32
Known via "static", distance 1, metric 0 (connected)
Tag 10
Redistributing via bgp 10
Advertised by bgp 10 route-map BLACKHOLE
Routing Descriptor Blocks:
* directly connected, via Null0
Route metric is 0, traffic share count is 1
Route tag 10

R3(config)# ip route 192.0.2.1 255.255.255.255 null0
R3(config)# do show ip bgp
BGP table version is 14, local router ID is 120.100.3.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
r RIB-failure, S Stale
```

```

Origin codes: i - IGP, e - EGP, ? - incomplete
Network          Next Hop          Metric LocPrf Weight Path
*>i126.1.1.0/24   120.100.1.1        0      100     0 i
*>i130.1.1.0/24   120.100.2.1        0      100     0 i
*>i130.100.200.0/24 120.100.2.1        0      100     0 i
* i150.100.2.100/32 192.0.2.1          0      100     0 i

```

R1(config)# **ip route 192.0.2.1 255.255.255.255 null0**

R1(config)# **interface Gigabit0/1**

R1(config-if)# **no icmp unreachable**

R1(config-if)# **do show ip bgp**

BGP table version is 8, local router ID is 126.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
r RIB-failure, S Stale

```

Origin codes: i - IGP, e - EGP, ? - incomplete
Network          Next Hop          Metric LocPrf Weight Path
*> 126.1.1.0/24   0.0.0.0           0           32768 i
*>i130.1.1.0/24   120.100.2.1        0      100     0 i
*>i130.100.200.0/24 120.100.2.1        0      100     0 i
* i150.100.2.100/32 192.0.2.1          0      100     0 i

```

R1# **show ip route 150.100.2.100**

Routing entry for 150.100.2.100/32
Known via "bgp 10", distance 200, metric 0, type internal
Last update from 192.0.2.1 00:00:02 ago
Routing Descriptor Blocks:
* 192.0.2.1, from 120.100.3.1, 00:00:02 ago
Route metric is 0, traffic share count is 1
AS Hops 0

R1# **show ip route 192.0.2.1**

Routing entry for 192.0.2.1/32
Known via "static", distance 1, metric 0 (connected)
Routing Descriptor Blocks:
* directly connected, via Null0
Route metric is 0, traffic share count is 1

- In a view of protecting the control plane on Router R6, configure CoPP so that IP Packets with a TTL of 0 or 1 are dropped rather than processed with a resulting ICMP redirect sent to the source. (1 point)

Cisco IOS Software sends all packets with a TTL of 0 or 1 to the process level to be processed. The device must then send an ICMP TTL expire message to the source. By filtering packets that have a TTL of 0 and 1, you can reduce the load on the process level. The control plane policing simply blocks packets with a TTL value of 0 and 1 as directed, but this will break your EIGRP and BGP peering. So you must specifically permit these packets within your ACL; otherwise, you would have just lost valuable points. If you found yourself running short on time and couldn't justify further time to investigate how to maintain your routing peering, remember that this is a 1-point question, worth leaving and coming back to if possible. If you have configured this correctly, as shown in Example 1-42, you have scored 1 point.

EXAMPLE 1-42 CoPP Configuration

```
R6(config)# ip access-list extended TTL
R6(config-ext-nacl)# deny eigrp any any
R6(config-ext-nacl)# deny tcp any any eq bgp
R6(config-ext-nacl)# deny tcp any eq bgp any
R6(config-ext-nacl)# permit ip any any ttl eq 0 1
R6(config-ext-nacl)# class-map DROP-TTL-0/1
R6(config-cmap)# match access-group name TTL
R6(config-cmap)# policy-map CoPP-TTL
R6(config-pmap)# class DROP-TTL-0/1
R6(config-pmap-c)# drop
R6(config-pmap-c)# control-plane
R6(config-cp)# service-policy input CoPP-TTL
```

Section 7: Multicast (4 Points)

- Configure Routers R1, R2, R3, and R4 for IPv4 Multicast. Configure R3 to send multicast advertisements of its own time by use of NTP sourced from interface Gig 0/0. Configure PIM sparse mode on all required interfaces. R3 should also be used to advertise its own gigabit interface IP address as an RP. R3 should also advertise the IP address you are using for the NTP advertisements, which will be 224.0.1.1. Do not use the command **ntp server** in any configurations. Routers R1, R2, and R4 should all show a clock synchronized to that of R3. (4 points)

NTP can be multicast on the reserved group IP address of 224.0.1.1 rather than the more familiar broadcast or unicast scenarios. The question requires you to configure R3 to become the NTP master and announce the group address to the NTP clients. As you are not permitted to use the command **ntp server** you must configure the clients with the command **ntp multicast client**. They will then have the capability to join the NTP group by use of PIM. It is good practice to TTL

scope your multicast announcements so that they do not propagate past the domain you require. If you haven't taken this into consideration in your solution, you would not be deducted points, but be aware of the facility in case you are met with a question that specifies this. If you have configured this correctly, as shown in Example 1-43, you have scored 4 points.

EXAMPLE 1-43 NTP Multicast Configuration and Verification

```
R3(config)# ip multicast-routing
R3(config)# ntp master
R3(config)# interface GigabitEthernet0/0
R3(config-if)# ip pim sparse-mode
R3(config-if)# ntp multicast ttl 2
R3(config-if)# interface Serial0/0/0
R3(config-if)# ip pim sparse-mode
R3(config-if)# ip pim send-rp-announce GigabitEthernet0/0 scope 2 group-list 4
R3(config)# ip pim send-rp-discovery GigabitEthernet0/0 scope 2
R3(config)# access-list 4 permit 224.0.1.1

R3# show ntp status
Clock is synchronized, stratum 8, reference is 127.127.7.1
nominal freq is 250.0000 Hz, actual freq is 250.0000 Hz, precision is
2**18 reference time is C98F1E61.2AE19310 (21:17:21.167 UTC Tue Feb 27
2007)
clock offset is 0.0000 msec, root delay is 0.00 msec
root dispersion is 0.02 msec, peer dispersion is 0.02 msec

R1(config)# ip multicast-routing
R1(config-if)# interface Serial0/0/0
R1(config-if)# ip pim sparse-mode
R1(config-if)# ntp multicast client

R1# show ntp status
Clock is synchronized, stratum 9, reference is 120.100.34.3
nominal freq is 250.0000 Hz, actual freq is 250.0000 Hz, precision is
2**18 reference time is C98F1E79.9FB2321D (21:17:45.623 UTC Tue Feb 27
2007)
clock offset is 0.0157 msec, root delay is 3.88 msec
root dispersion is 0.06 msec, peer dispersion is 0.02 msec
R1(config-if)#
R1# show ip igmp group
IGMP Connected Group Membership
```

Group Address	Interface	Uptime	Expires	Last Reporter
224.0.1.1	Serial0/0/0	00:40:12	00:02:50	120.100.123.1
224.0.1.39	Serial0/0/0	00:07:21	00:02:51	120.100.123.3
224.0.1.40	Serial0/0/0	00:40:13	00:02:52	120.100.123.1

```
R2(config)# ip multicast-routing
R2(config-if)# interface Serial0/0
R2(config-if)# ip pim sparse-mode
R2(config-if)# ntp multicast client
```

R2# show ntp status

```
Clock is synchronized, stratum 9, reference is 120.100.34.3
nominal freq is 250.0000 Hz, actual freq is 250.0000 Hz, precision is 2**18 reference time is C98F1E73.83B73E68 (21:17:39.514 UTC Tue Feb 27 2007)
clock offset is 0.0182 msec, root delay is 4.14 msec
R2# show ip igmp group
```

```
root dispersion is 15875.06 msec, peer dispersion is 15875.02 msec
```

IGMP Connected Group Membership

Group Address	Interface	Uptime	Expires	Last Reporter
224.0.1.1	Serial0/0	00:41:08	00:02:59	120.100.123.2
224.0.1.40	Serial0/0	a		
a00e:41:d09	Serial0/0	00:01:59	120.100.123.2	
224.0.1.39	Serial0/0	00:08:12	00:02:57	120.100.123.3

```
R4(config)# ip multicast-routing
R4(config-if)# interface GigabitEthernet0/0
R4(config-if)# ip pim sparse-mode
R4(config-if)# ntp multicast client
```

R4# show ntp status

```
Clock is synchronized, stratum 9, reference is 120.100.34.3
nominal freq is 250.0000 Hz, actual freq is 250.0000 Hz, precision is
2**18 reference time is C98F1EF1.2B7DB1F2 (21:19:45.169 UTC Tue Feb 27
2007)
clock offset is -0.6937 msec, root delay is 1.37 msec
root dispersion is 7877.08 msec, peer dispersion is 7876.34 msec
```

R4# show ip igmp group

IGMP Connected Group Membership

Group Address	Interface	Uptime	Expires	Last Reporter
224.0.1.1	GigabitEthernet0/0	00:41:29	00:02:42	120.100.34.4

224.0.1.39	GigabitEthernet0/0	00:08:35	00:02:42	120.100.34.3
224.0.1.40	GigabitEthernet0/0	00:41:07	00:02:42	120.100.34.4

IP Services (4 Points)

- Configure the following commands on Router R1:

```
aaa new-model

logging buffered

logging 120.100.99.1
```

Configure a policy on Router R1 so that if a user tries to remove AAA services or disable logging via the CLI that a syslog message of UNAUTHORIZED-COMMAND-ENTERED is generated. The policy should ensure either command is not executed and should consist of a single-line command for the CLI pattern detection. The policy and CLI should run asynchronously. The policy should also generate an email from the router to a mail server residing on IP address 120.100.99.2 (to security@lab-exam.net from eem@lab-exam.net, with the subject "User-Issue," with the message body consisting of details of who was logged on the time either of the commands were entered). (2 points)

This is an intricate Embedded Events Manager (EEM) question. You are required to configure an EEM applet with a CLI pattern event on a single line to match on either of the commands (**no aaa xxx** and **no logging xxx**). This is achieved by a pattern of `^no (aaa|logging).*`. The following **sync no skip yes** parameters simply state that the policy and CLI should run asynchronously and that the command entered should not be executed as directed. When the commands are matched via the CLI pattern, the policy requires the syslog message to be generated, a CLI command action to run "show users," and a final action to send an email with the details of the previous **show** command (which is achieved by the command `$_cli_result`). Example 1-44 details the required configuration and resulting execution of the EEM when the commands **no aaa new-model** and **no logging buffered** are entered and not executed on the router. If you have configured this correctly, as shown in Example 1-44, you have scored 2 points.

EXAMPLE 1-44 R1 EEM Configuration and Verification Testing

```
R1(config)# aaa new-model
R1(config)# logging buffered
R1(config)# logging 120.100.99.1
R1(config)#
R1(config)# event manager applet CCIE-QUESTION
```

```
R1(config-applet)# event cli pattern "^no (aaa|logging).*" sync no skip yes
R1(config-applet)# action 1.0 syslog msg "UNAUTHORIZED-COMMAND-ENTERED"
R1(config-applet)# action 2.0 cli command "show user"
R1(config-applet)# action 3.0 mail server "120.100.99.2" to "security@lab-exam.net" from "eem@lab-exam.net" subject "User-Issue" body "$_cli_result"

R1(config-applet)# no aaa new-model
%HA_EM-6-LOG: CCIE-QUESTION: UNAUTHORISED-COMMAND-ENTERED
%HA_EM-3-FMPD_SMTP_CONNECT: Unable to connect to SMTP server: 120.100.99.2
%HA_EM-3-FMPD_ERROR: Error executing applet CCIE-QUESTION statement 3.0
R1(config)# no logging buffered
%HA_EM-6-LOG: CCIE-QUESTION: UNAUTHORISED-COMMAND-ENTERED
%HA_EM-3-FMPD_SMTP_CONNECT: Unable to connect to SMTP server: 120.100.99.2
%HA_EM-3-FMPD_ERROR: Error executing applet CCIE-QUESTION statement 3.0
R1(config)# do show run | include aaa new-model
aaa new-model
R1(config)# do show run | include logging buffered
logging buffered 4096 debugging
```

- Cisco WAAS devices are to be installed on Switches 1 and 2 in the future on VLAN300. Configure Routers R5 and R6 to provide WCCPv2 redirection for clients residing on VLAN300 to ensure that all TCP traffic other than telnet is redirected only to the WAEs, which will reside on addresses 150.100.3.50 and .51 within VLAN300. You are not required to configure the switches for WCCP and can assume that incoming WAAS traffic from the network will arrive at interfaces Gi0/0 on both R5 and R6. Secure your WCCP with a password of CCIE. (2 points)

WCCP in this scenario could be configured on the routers or Switches 1 and 2, but you are directed to configure the routers. WCCP service 62 is used to redirect traffic sourced on VLAN300, which is applied to the VLAN300 interfaces of R5 and R6. And WCCP service 61 is used for the redirection of the incoming traffic, which is applied as directed to Gi0/0 on both R5 and R6. Telnet traffic is excluded (generally, management traffic is not recommended to be optimized) by creation of an extended ACL, which is applied to services 61 and 62 in a redirect-list. You need to remember to permit all other TCP and not just IP because the WAE can optimize only TCP sessions. The WAE devices are included in a group-list for services 61 and 62, and a password is applied as directed. (The group-list will aid in load sharing and can stop a bogus WCCP device from attempting to receive redirected traffic.) If you have configured this correctly, as shown in Example 1-45, you have scored 2 points.

EXAMPLE 1-45 R5 and R6 WCCP Configuration

```
R5(config)# ip wccp 61 password 0 CCIE
R5(config)# ip wccp 62 password 0 CCIE
R5(config)# ip access-list extended WAAS
R5(config-ext-nacl)# remark WAAS DENY TELNET
R5(config-ext-nacl)# deny tcp any any eq telnet
R5(config-ext-nacl)# deny tcp any eq telnet any
R5(config-ext-nacl)# permit tcp any any
R5(config-ext-nacl)# exit
R5(config)# ip wccp 61 group-list 2 redirect-list WAAS
R5(config)# ip wccp 62 group-list 2 redirect-list WAAS
R5(config)# access-list 2 permit 150.100.3.50
R5(config)# access-list 2 permit 150.100.3.51
R5(config)# interface Gi0/1
R5(config-if)# ip wccp 61 redirect in
R5(config-if)# interface Gi0/0
R5(config-if)# ip wccp 62 redirect in

R6(config)# ip access-list extended WAAS
R6(config-ext-nacl)# remark WAAS DENY TELNET
R6(config-ext-nacl)# deny tcp any any eq telnet
R6(config-ext-nacl)# deny tcp any eq telnet any
R6(config-ext-nacl)# permit tcp any any
R6(config-ext-nacl)# exit
R6(config)# ip wccp 61 group-list 2 redirect-list WAAS
R6(config)# ip wccp 62 group-list 2 redirect-list WAAS
R6(config)# access-list 2 permit 150.100.3.50
R6(config)# access-list 2 permit 150.100.3.51
R6(config)# interface Gi0/1
R6(config-if)# ip wccp 61 redirect in
R6(config-if)# # interface Gi0/0
R6(config-if)# ip wccp 62 redirect in
```

Lab WRAP-UP

So how did it go, did you run out of time, did you manage to finish but miss what was actually required? If you scored over 80, then well done. If you accomplished this within the time frame of 8 hours or less, you will be prepared for any scenario that you are likely to face during the 5 1/2 hours of the Configuration section of the actual exam. Remember that the troubleshooting section on the v4.0 exam is a separate section than the configuration with a different scenario, and you will have 2 hours to complete this. This lab was designed to ensure you troubleshoot your own work as you progress through the questions. What sets the CCIE exam apart within the industry is the complexity of the questions to test you further than you thought possible. The exam isn't trying to trick you, but it will ensure that you have the ability to think laterally—an ability that will ensure you exceed in your networking career and one that sets CCIEs apart. Spend the time to go back over the questions and practice with the configurations using **debug** and **show** commands to fully absorb any new areas you might have come across.

Did you anticipate and factor into your configuration items such as the offset-list within RIPng for the tunnel and maximum reserved bandwidth within QoS? If you did, congratulations, because this would have saved you time and secured you points. It also shows that you fully understand the protocols involved and adapt at testing your configurations. How can you ensure that you have the ability to spot any underlying issues related to a question? Well it's all mileage; you'll get out of your study what you put into it.

Practice Lab 2

The CCIE exam commences with 2 hours of troubleshooting followed by 5 1/2 hours of configuration and a final 30 minutes of additional questions. This lab has been timed to last for 8 hours of configuration and self-troubleshooting, so aim to complete the lab within this period.

Then either score yourself at this point or continue until you feel you have met all the objectives. You now are going to be guided through the equipment requirements and pre-lab tasks in preparation for taking this practice lab.

If you don't own six routers and four switches, consider using the equipment available and additional lab exercises and training facilities available within the CCIE R&S 360 program. Detailed information on the 360 program and CCIE R&S exam can be found on the following URLs, respectively:

https://learningnetwork.cisco.com/community/learning_center/cisco_360/360-rs

https://learningnetwork.cisco.com/community/certifications/ccie_routing_switching

NOTE

The 3825s used in this lab were loaded with **c3825-adventerprisek9-mz.124-6.T.bin**, and the 3725 was loaded with **c3725-adventerprisek9-mz.124-6.T.bin**.

Equipment List

You will need the following hardware and software components to begin this practice lab.

- Six routers loaded with Cisco IOS Software Release 12.4 Advanced Enterprise image and the minimum interface configuration as documented in Table 2-1

TABLE 2-1 Hardware Required per Router

Router	Model	Ethernet I/F	Serial I/F
R1	3825	1	1
R2	3725	1	1
R3	3825	1	1
R4	3825	2	—
R5	3825	2	—
R6	3825	2	—

NOTE

The 3550 in this lab was loaded with **c3550-ipservicesk9-mz.122-25.SEE.bin**, and the 3560s with **c3560-ipservicesk9-mz.122-25.SEE.bin**.

NOTE

Notice in the initial configurations supplied that some interfaces will not have IP addresses pre-configured. This is because you will either not be using that interface or you must configure it from default within the exercise. The initial configurations supplied should be used to pre-configure your routers and switch before the lab starts.

If your routers have different interface speeds than those used in this book, adjust the bandwidth statements on the relevant interfaces to keep all interface speeds in line. This will ensure that you do not get unwanted behavior because of differing IGP metrics.

- One 3550 switch with IOS 12.2 IP Services and 3 3560 Switches with IOS 12.2 IP Services.

Setting Up the Lab

Use any combination of routers as long as you fulfill the requirements within the topology diagram, as shown in Figure 2-1. However, it is recommended to use the same model of routers because this can make life easier if you load configurations directly from the supplied configurations into your own devices.

Lab Topology

This practice Lab uses the topology as outlined in Figure 2-1, which you will need to re-create with your own equipment or by using lab equipment on the CCIE R&S 360 program.

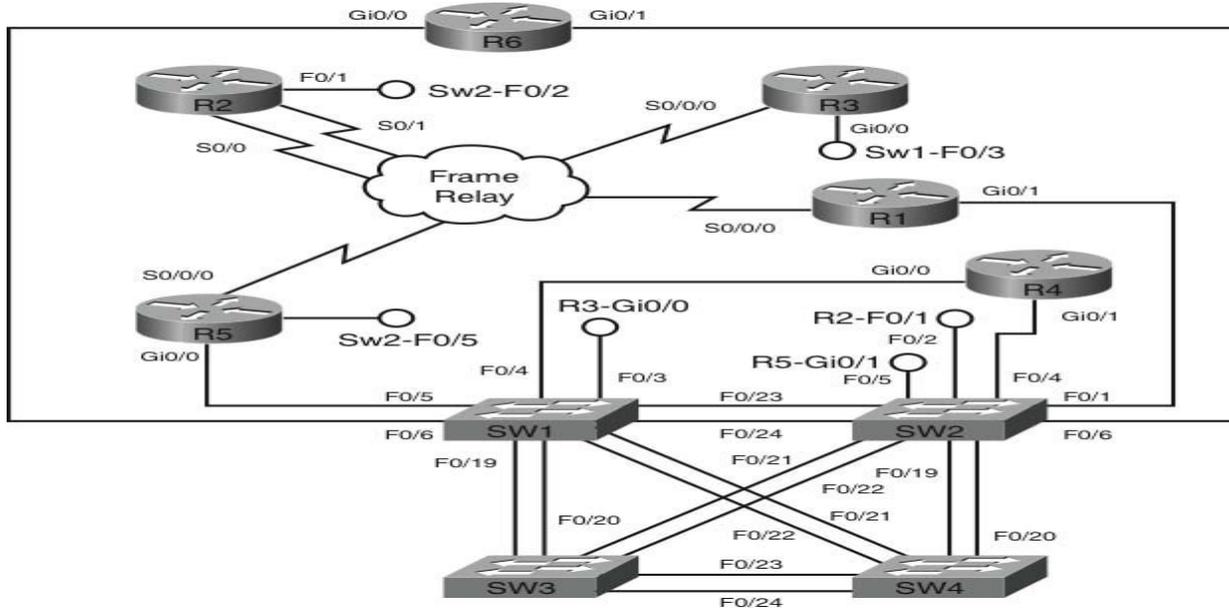


FIGURE 2-1
Lab 1 Topology
Diagram

Switch Instructions

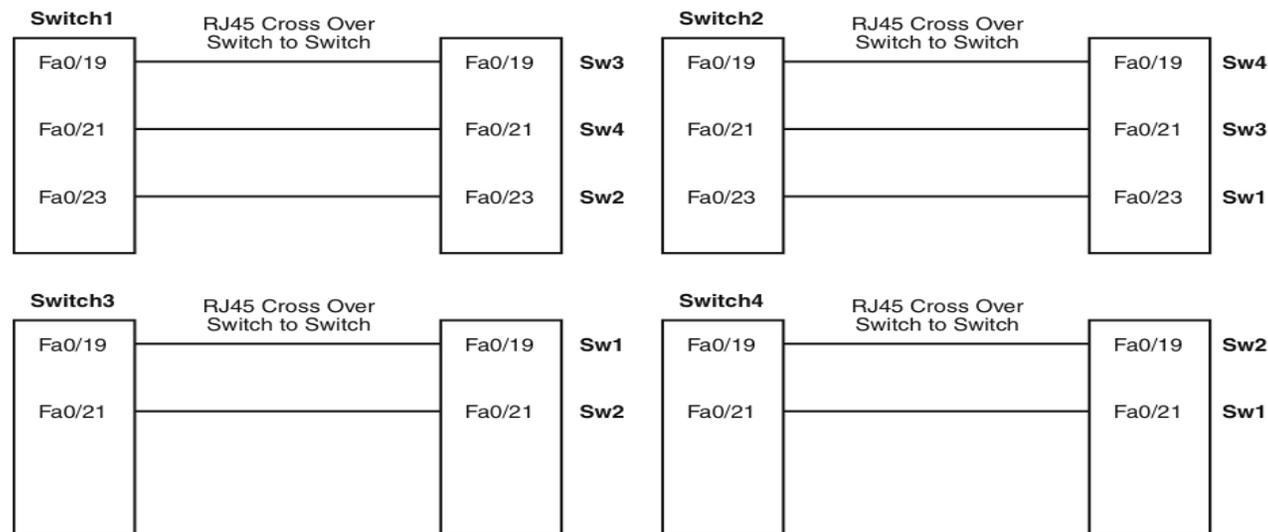
Configure VLAN assignments from the configurations supplied or from Table 2-2.

TABLE 2-2 VLAN Assignment

VLAN	Switch1	Switch2	Switch3	Switch4
34	Fa0/3, Fa0/4, Fa0/5	—	—	—
46	Fa0/6	Fa0/4	—	—
53	VLAN53	Fa0/5	VLAN53	—
63	—	Fa0/6	VLAN63	VLAN63
100	—	Fa0/1	—	—
200	—	Fa0/2	—	—

Connect your switches with RJ45 Ethernet cross over cables, as shown in Figure 2-2.

FIGURE 2-2
Switch to Switch
Connectivity



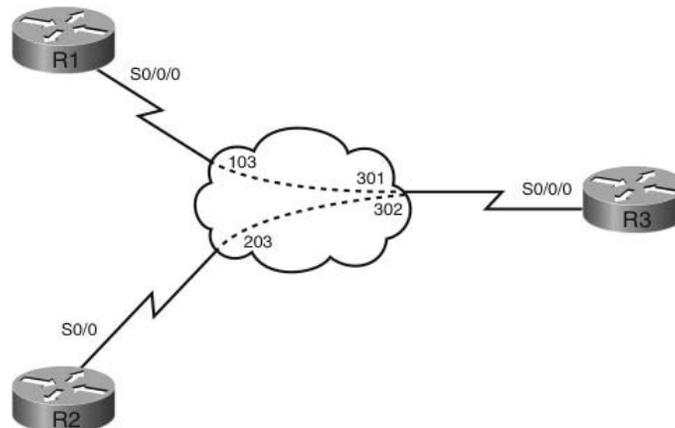
Frame Relay Instructions

Configure one of the routers you are going to use in the lab as a Frame Relay switch, or have a dedicated router purely for this task. This lab uses a dedicated router for the Frame Relay switch. A fully meshed environment is configured between all the Frame Relay routers. Pay attention in the lab as to which permanent virtual circuits (PVC) are actually required. Keep the encapsulation and Local Management Interface (LMI) settings to default for this exercise, but experiment with the settings outside these labs because you could be required to configure the Frame Relay switching within your real lab.

If you are using your own equipment, keep the DCE cables at the frame switch end for simplicity and provide a clock rate to all links from this end.

After configuration, the Frame Relay connectivity will represent the logical Frame Relay network, as shown in Figure 2-3.

FIGURE 2-3
Frame Relay Logical
Connectivity

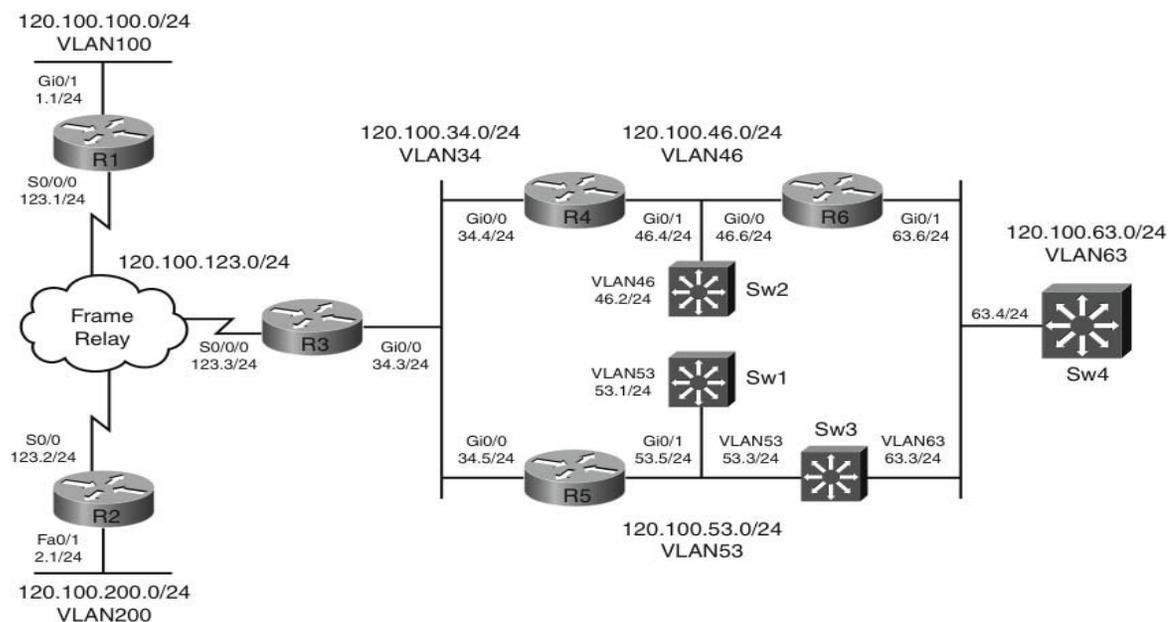


IP Address Instructions

You will find in the actual CCIE lab that the majority of your IP addresses will be preconfigured. For this exercise you are required to configure your IP addresses, as shown in Figure 2-4, or load the initial router configurations supplied. If you are manually configuring your equipment, ensure you include the following Loopback addresses (R1 and R3 use the same IP address for Loopback 255):

R1 Lo0 120.100.1.1/24	R6 Lo0 120.100.6.1/24
Lo255 200.200.200.200/24	SW1 Lo0 120.100.7.1/24
R2 Lo0 120.100.2.1/24	SW2 Lo0 120.100.8.1/24
R3 Lo0 120.100.3.1/24	SW3 Lo0 120.100.9.1/24
Lo255 200.200.200.200/24	SW4 Lo0 120.100.10.1/24
R4 Lo0 120.100.4.1/24	
R5 Lo0 120.100.5.1/24	

FIGURE 2-4
IP Addressing Diagram



Pre-lab Tasks

- Build the lab topology per Figure 2-1 and Figure 2-2.
- Configure your Frame Relay switch router to provide the necessary Data Link Control Identifiers (DLCI) per Figure 2-3.
- Configure the IP addresses on each router as shown in Figure 2-4 and add the Loopback addresses. Alternatively, you can load the initial configuration files supplied if your router is compatible with those used to create this exercise.

General Guidelines

- Read the whole lab before you start.
- Do not configure any static/default routes unless otherwise specified.

- Use only the DLCIs provided in the appropriate figures.
- Ensure full IP visibility between routers for ping testing/telnet access to your devices.
- If you run out of time, choose questions that you are confident you can answer, or choose questions with a higher point rating to maximize your potential score.
- Get into a comfortable and quiet environment where you can focus for the next 8 hours.
- Take a 30-minute break midway through the exercise.
- Have available a Cisco Documentation CD-ROM or access online the latest documentation from the following URLs:

www.cisco.com/univercd/home/home.htm

http://www.cisco.com/en/US/products/ps6350/products_installation_and_configuration_guides_list.html

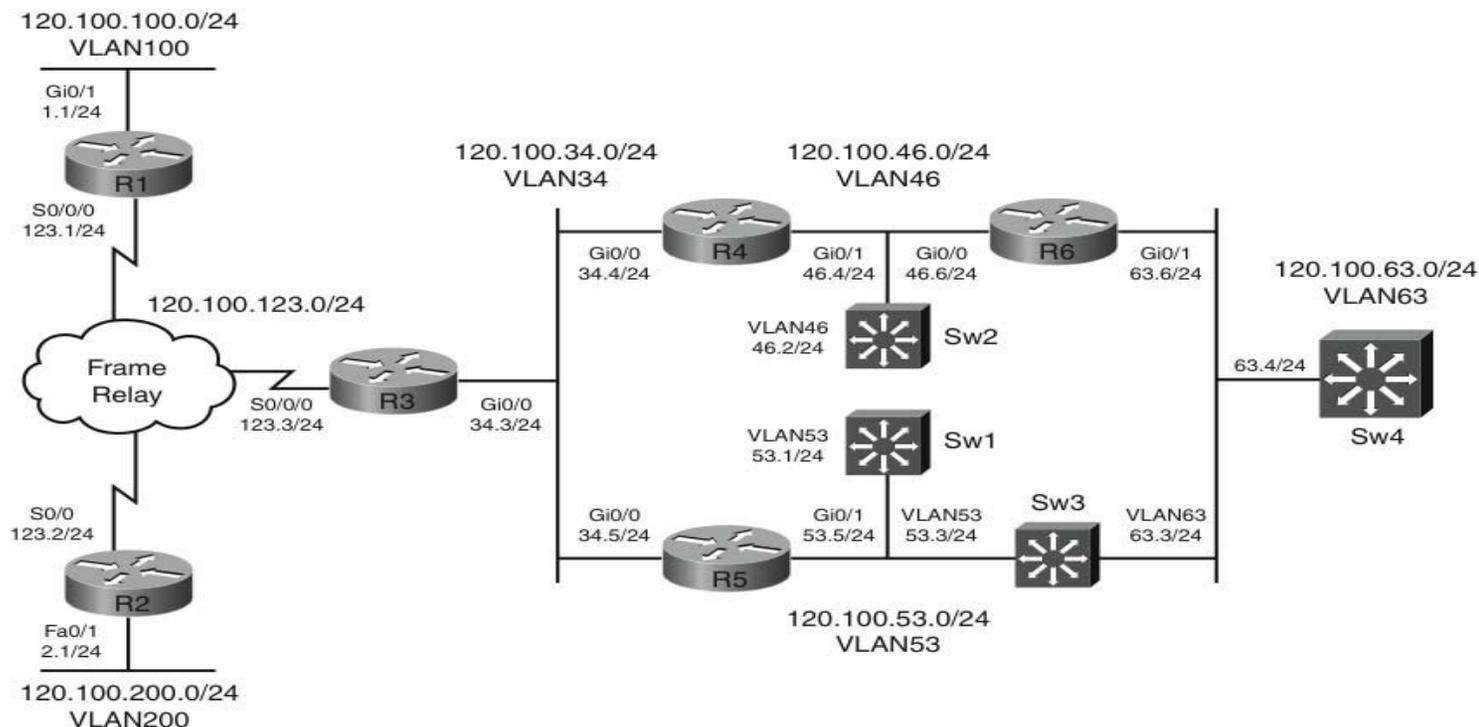
NOTE

Access only these URLs, not the whole Cisco.com website because if you are permitted to use documentation during your CCIE lab exam, it will be restricted. Consider opening several windows with the pages you are likely to look at to save time during your lab.

Practice Lab Two

You will now be answering questions in relation to the network topology, as shown in Figure 2-5.

FIGURE 2-5
Lab Topology Diagram

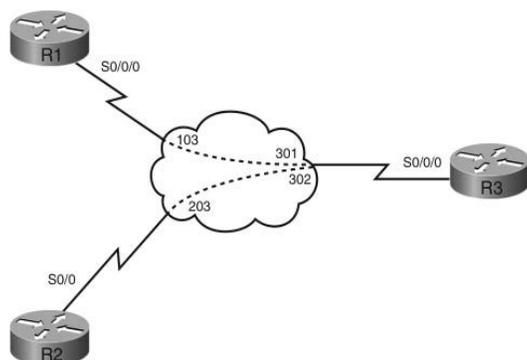


Section 1: LAN Switching and Frame-Relay (24 Points)

- Configure your switched network to use 802.1w Spanning Tree. Switch 1 should be the root bridge for VLANs 34,46,53,63,100 and 200, with Switch 2 being the secondary root bridge for all listed VLANs. (3 points)
- Switch 3 should use its interface directly connecting to Switch 2 (Fast Ethernet 0/21) for traffic directed toward even-numbered VLANs (34, 46, 100, 200) and the interface directly connecting to Switch 1 (Fast Ethernet 0/19) for odd-numbered VLANs (53, 63). (3 points)
- Switch 4 should use its interface directly connecting to Switch 2 (Fast Ethernet0/19) for traffic destined toward even-numbered VLANs (34, 46, 100, 200) and the interface directly connected to Switch 1 (Fast Ethernet 0/21) for odd-numbered VLANs (53, 63). (3 points)
- Ensure a cable fault between Switches 1 and 2 could not result in one-way traffic between the two switches, resulting in spanning-tree issues. (2 points)

- Configure Switch 1 and Switch 2 to enable connectivity of two further switches in the future to be connected to ports Fast Ethernet 0/18 on each switch. The new switches should be able to tunnel their own configured VLANs through a new VLAN (30) between Switch 1 and Switch 2. There is no requirement to configure a root bridge or VLAN load balancing for the new VLAN between Switch 1 and Switch 2. (4 points)
- Configure your switched network to monitor the VLAN200 interface associated with R2 (Switch 2 FastEthernet 0/1), and send only traffic destined to R2 on this switch port across your network to Switch 3 port Fast Ethernet 0/17—use a new VLAN (20) to assist in this configuration. There is no requirement to configure a root bridge or VLAN load balancing for the new VLAN. (3 points)
- Configure the interface on Switch 2 that connects to R5 VLAN53 (Fast Ethernet 0/5) in such a way that if all the trunks on Switch 2 connecting to Switch 1, Switch 3, and Switch 4 should fail, this Ethernet port transitions into error-disable state. (3 points)
- Configure interfaces Fast Ethernet 0/9 and 0/10 on Switch 1 so that even if they are configured to belong to the same VLAN they will not be able to forward unicast, broadcast, or multicast traffic to one another. Do not use any form of ACL or configure the ports to belong to a PVLAN. (1 point)
- Your initial Frame-Relay configuration has been supplied for the R1-R2-R3 connectivity. Configure Frame Relay per Figure 2-7 to ensure each device is reachable over the Frame-Relay network. Use only the indicated DLCIs, and ensure that a proprietary method of reducing the payload over the Frame-Relay network is enabled on a per-packet basis. (2 points)

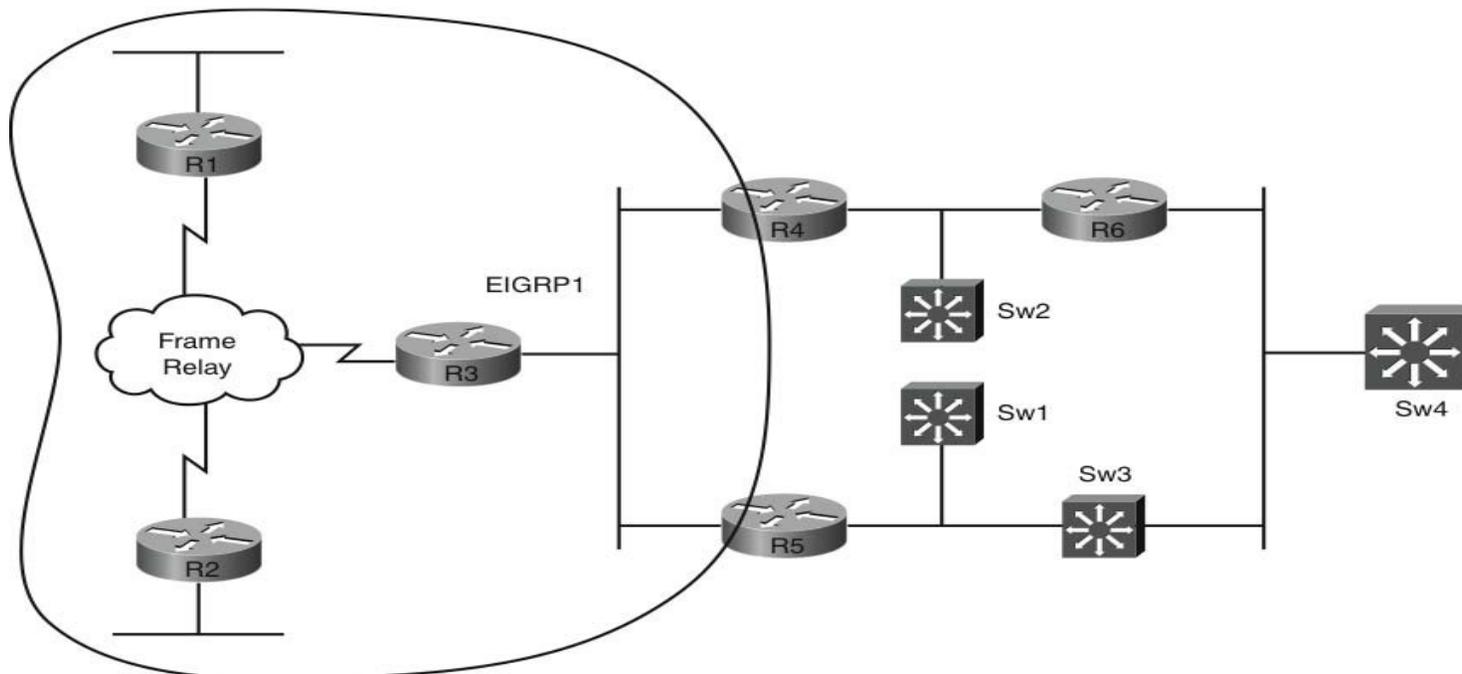
FIGURE 2-6
Frame Relay Connectivity



Section 2: IPv4 IGP Protocols (28 Points)

Section 2.1: EIGRP

FIGURE 2-7
EIGRP Topology

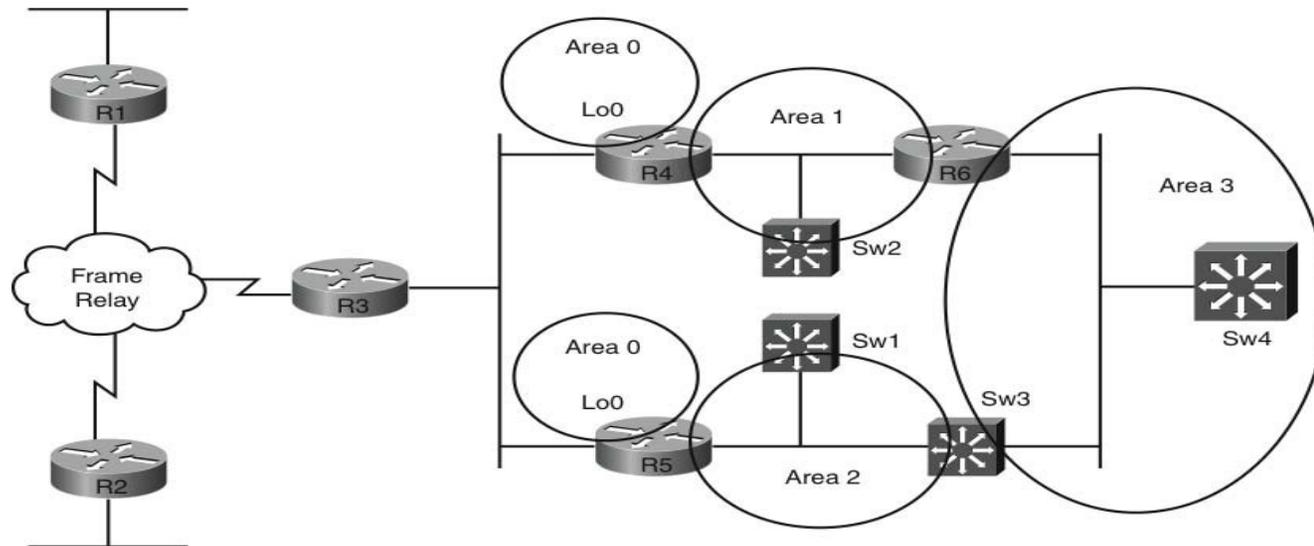


- Configure EIGRP per Figure 2-7 using an AS of 1; each EIGRP router should have its Loopback 0 interface configured and advertised within EIGRP. (2 points)
- Configure R1 to advertise a summary route of 120.100.0.0/16 outbound on its serial interface. R3 should see the original VLAN100 and Loopback 0 individual routes in addition to the summary route. You may use only one summary route in your configuration. (3 points)
- Ensure the length of time that EIGRP considers neighbors to be valid without receiving a hello packet on the Frame-Relay network between R1, R2, and R3 is 200 seconds; do not change the hello-interval parameter. (2 points)

- Configure new Loopback interfaces on R1 and R2 using a Loopback interface 2 with an identical IP address of 150.101.1.1/24 on both routers; advertise this network into EIGRP on each router. Ensure that R3 prefers the route from R2 by manipulating the delay associated with this route. Do not manually adjust the delay associated with the interface by use of the `delay` command. You are only permitted to configure R2 to influence the delay. (3 points)

Section 2.2: OSPF

FIGURE 2-8
OSPF Topology



- Configure OSPF per Figure 2-8 using a process ID of 1. All OSPF configuration, where possible, should not be configured under the process ID. Each OSPF router should also have its Loopback 0 interface configured and advertised within OSPF as follows: (2 points)
 - R4 Loopback 0 – Area 0
 - R5 Loopback 0 – Area 0
 - R6 Loopback 0 – Area 1
 - Sw1 Loopback 0 – Area 2

Sw2 Loopback 0 – Area 1

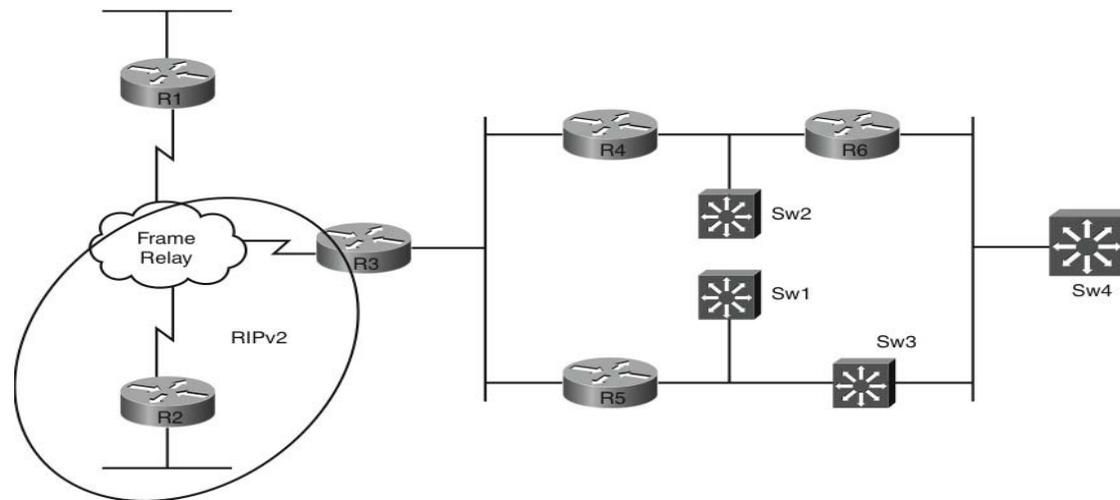
Sw3 Loopback 0 – Area 2

Sw4 Loopback 0 – Area

- Area 0 is partitioned between R4 and R5—ensure your network can accommodate this issue. You are not permitted to form any area 0 neighbor relationship directly between R4 and R5 to join area 0. (4 points)

Section 2.2: RIPv2

FIGURE 2-9
RIPv2 Topology



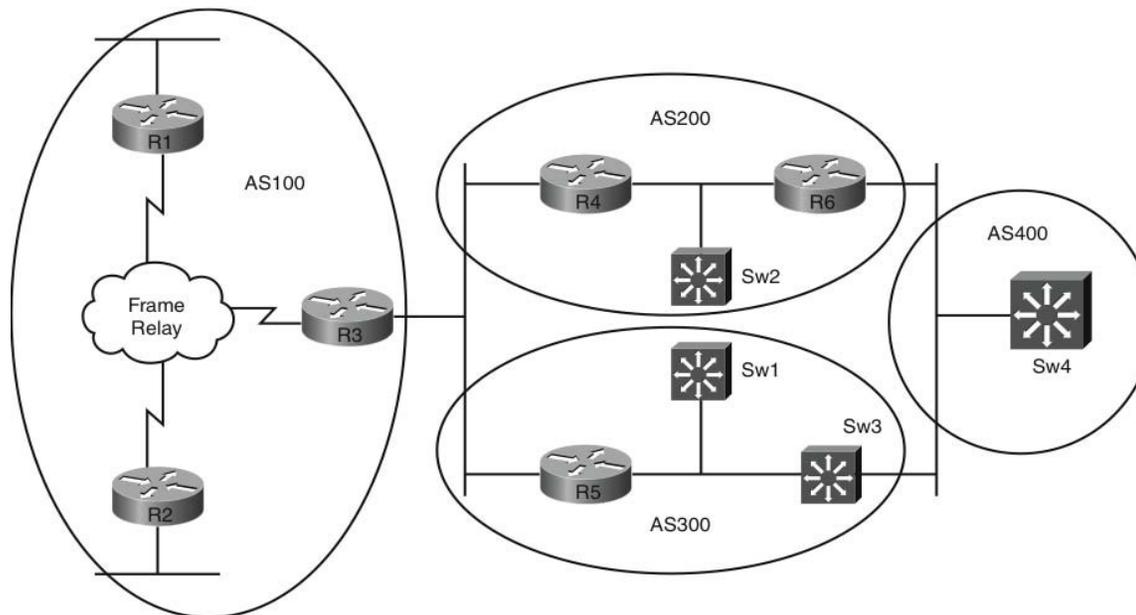
- Configure RIPv2 between R2 and R3, configure a new Loopback interface on R2 (Loopback 3) with an IP address of 150.101.2.1/24, and advertise this and only this network to R3 from R2. (2 points)
- R3 should not advertise any connected interfaces into RIPv2; do not filter routing advertisements to achieve this behavior. (2 points)

Section 2.3: Redistribution

- Perform a one-way redistribution of RIPv2 into EIGRP on R3 using the following default metric: 1544 20000 255 1 1500. Ensure that R1 shows a next hop for the RIPv2 advertised route of 150.101.2.0/24 of R2 and perform configuration only on R3 for this task. (3 points)
- Perform mutual redistribution of EIGRP and OSPF on R4 and R5. Use a metric of 5000 for redistributed routes into OSPF that should appear as external type 2 routes and the following K values for OSPF routes redistributed into EIGRP: 1544 20000 255 1 1500. (2 points)
- R3 will have equal cost external EIGRP routes to the redistributed OSPF subnet 120.100.63.0/24 (VLAN 63). Configure only R3 to ensure that R3 routes via a next hop of R5 (120.100.34.5) for this destination subnet. If this route fails, the route advertised from R4 (120.100.34.4) should be used dynamically. (3 points)

Section 3: BGP (15 Points)

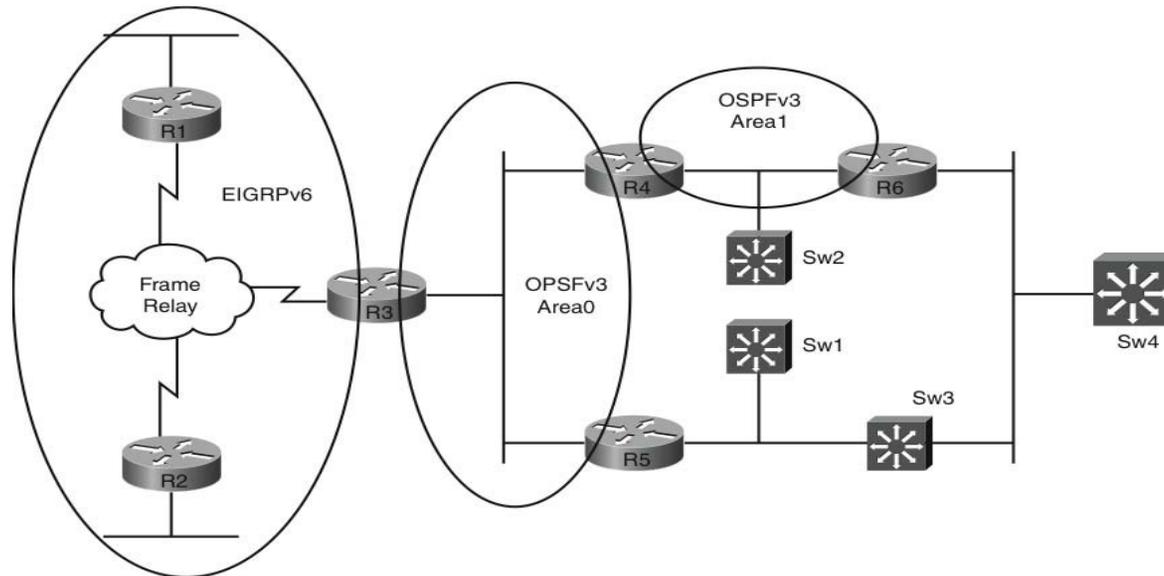
FIGURE 2-10
BGP Topology



- Configure BGP peering per Figure 2-9 as follows: iBGP R1-R3, R2-R3, R4-R6, R4-SW2. R5-Sw1 R5-sw3. eBGP R3-R4, R3-R5, Sw4-Sw3. R6-Sw4. Use Loopback interfaces to peer on all routers with the exception of peering between R3-R4 and R3-R5. Do not use the command **ebgp-multihop** within your configurations. (3 points)
- Routers R1 and R2 in AS100 should be made to only passively accept BGP sessions. R3 should be configured to only actively create BGP sessions to R1 and R2 within AS100. (3 points)
- Configure the following Loopback interfaces on R3 and Sw4; advertise these networks into BGP using the **network** command: (2 points)
 - R3 – Loopback interface 5 (152.100.100.1/24)
 - Sw4 – Loopback interface 5 (152.200.32.1/24)
 - Sw4 – Loopback interface 6 (152.200.33.1/24)
 - Sw4 – Loopback interface 7 (152.200.34.1/24)
 - Sw4 – Loopback interface 8 (152.200.35.1/24)
- Configure R3 to inform R4 that it does not want to receive routes advertised from Sw4 for networks 152.200.33.0/24, 152.200.34.0/24, and 152.200.35.0/24. Achieve this in such a manner that R4 does not actually advertise these routes toward R3. You may also configure R4. (4 points)
- Configure a route-map on R5 that prepends its local AS 2 an additional two times for network 152.200.32.0/24 when advertised to R3. The route-map may contain multiple permit statements but only one prepend is permitted per line. (3 points)

Section 4: IPv6 (12 Points)

FIGURE 2-11
IPv6 Topology



Configure IPv6 addresses on your network as follows:

- 2007:C15:C0:10::1/64 - R1 Gi0/0
- 2007:C15:C0:11::1/64 – R1 tunnel0
- 2007:C15:C0:11::3/64 – R3 tunnel0
- 2007:C15:C0:12::2/64 - R2 tunnel0
- 2007:C15:C0:12::3/64 – R3 tunnel1
- 2007:C15:C0:13::2/64 – R2 fe0/1
- 2007:C15:C0:14::3/64 – R3 Gi0/0

2007:C15:C0:14::4/64 – R4 Gi0/0

2007:C15:C0:14::5/64 – R5 Gi0/0

2007:C15:C0:15::5/64 – R4 Gi0/1

2007:C15:C0:15::6/64 – R6 Gi0/0

Section 4.1: EIGRPv6

- Configure EIGRPv6 between R1, R2, and R3. EIGRPv6 should be enabled on the Ethernet interfaces of R1 and R2 and on all tunnel interfaces of R1, R2, and R3. Build your tunnels using **ipv6ip** mode—use an AS number of 6 on all required interfaces. (2 points)

Section 4.2: OSPFv3

- Configure OSPFv3 per Figure 2-11; use an OSPFv3 process of 1 on each router. (2 points)
- Configure Area 1 with IPsec authentication, use Message Digest 5, a security policy index of 500, and a key of **DECODECC1E0DDBA11B0BB0BBEDB00B00** (2 points)
- Ensure the area router in Area 1 receives the following route; you may configure R4 to achieve this: (2 points) OI

2007::/16 [110/2]

via XX.XX.XX.XX:XX.XX.XX.XX, GigabitEthernet0/0

Section 4.3: Redistribution

- Redistribute EIGRPv6 into OSPFv3 on R3. Redistributed EIGRPv6 routes should have a metric of 5000 associated with them, regardless of which area they are seen in within the OSPFv3 network. (2 points)
- Configure R3 so that both R1 and R2 have the following IPv6 EIGRPv6 route in place. Do not redistribute OSPF into EIGRPv6 to achieve this, and ensure all routers have full visibility. (2 points) D

2007::/16 [90/XXXXXXXXXX]

via XXXX::XXXX:XXXX:XXXX:XXXX, Tunnel0

Section 5: QoS (6 Points)

- 2 IP Video Conferencing units are to be installed onto Switch 2 ports FastEthernet 0/15 and 0/16 on VLAN 200. The devices use TCP ports 3230–3231 and UDP ports 3230–3235, and this traffic is unmarked from the devices as it enters the switch. Configure Switch 2 to assign a DSCP value of AF41 to video traffic from both of these devices. Ensure that the switch ports assigned to the devices do not participate in the usual spanning-tree checks, cannot form trunk links, and cannot be configured as Etherchannels. (3 points)
- Configure R2 to assign a strict priority queue with a 40 percent reservation of the WAN bandwidth for the Video Conferencing traffic in the previous question. Maximize the available bandwidth by ensuring the RTP headers within the video stream are compressed. The remainder of the bandwidth should be guaranteed for a default queue with WRED enabled. Assume the full line rate of 1.544 Mbps as the available WAN bandwidth, and ensure the complete bandwidth is utilized by both queues. (3 points)

Section 6: Multicast (7 Points)

- Configure Routers R1, R2, R3, and R4 for IPv4 multicast. Each router should use PIM sparse dense mode. Both R1 and R2 should be configured to be candidate RPs specifically for the following multicast groups: 225.225.0.1, 225.225.0.2, 225.225.0.3, and 225.225.0.4 by use of their Loopback 0 interfaces. You should limit the boundary of your multicast network so it does not propagate further into your network than R4. R3 should be configured as a mapping agent to announce the rendezvous points for the multicast network with the same boundary constraints. (3 points)
- Configure R3 to ensure R4 has a candidate RP as R1 for groups 225.225.0.1 and 225.225.0.2 and R2 for groups 225.225.0.3 and 225.225.0.4. (2 points)
- Configure R1 to monitor traffic forwarded through itself for traffic destined to the multicast group of 225.225.0.1. If no packet for this group is received within a single 10-second interval, ensure an SNMP trap is sent to an SNMP management station on 120.100.100.100 using a community string of “public.” (2 points)

Section 7: Security (7 Points)

- Allow Router R6 to passively watch the SYN connections that flow to only VLAN63 for servers that might reside on this subnet. To prevent a potential denial of service (DoS) attack from a flood of SYN requests, the router should be configured to randomly drop SYN packets from any source to this VLAN that have not been correctly established within 20 seconds. (2 points)

- Configure an ACL on R1 to allow TCP sessions generated on this router and through its Ethernet interface and to block TCP sessions from entering on its Frame-Relay interface that were not initiated on it or through it originally. Do not use the established feature within standard ACLs to achieve this, and apply ACLs only on the Frame-Relay interface. The ACL should timeout after 100 seconds of locally initiated TCP inactivity; it should also enable ICMP traffic inbound for testing purposes. (3 points)
- Configure R1 so it can perform SCP. The router should belong to a domain of toughest.co.uk; use local authentication with a username and password of cisco, a key size of 768 bits, and an SSH timeout of 2 minutes and retry value of 2. (2 points).

NOTE

This section should be used only if you require clues to complete the questions. In the actual CCIE lab, the proctor will not enter into any discussions regarding the questions or answers; he or she will be present to ensure you do not have problems with the lab environment and to maintain the timing element of the exam.

“Ask the Proctor”

Section 1: LAN Switching and Frame-Relay

Q: Do you just want me to configure the root and secondary root bridges into 802.1w spanning tree?

A: You should ensure that your network runs a consistent version of spanning tree.

Q: Can I change the root bridge assignments of odd- and even-numbered VLANs to ensure different interfaces are used on Switch 3 and Switch 4?

A: No, the root bridge assignment should remain as per the first question.

Q: If a copper Ethernet cable fails between Switch 1 and Switch 2, surely I wouldn't encounter spanning-tree issues because there would not be any loops present. Am I correct in thinking this?

A: Not entirely, consider a partial failure rather than a complete breakage.

Q: The switches are connected with Ethernet copper cables; wouldn't a feature like UDLD be beneficial only if the connections are fiber?

A: UDLD can operate over copper Ethernet in the same manner as Fiber.

Q: Would you like me to configure a native VLAN of 30 on trunks to the two new switches?

A: No, a native VLAN would not facilitate transportation of multiple VLANs over the single VLAN 30 between Switch 1 and Switch 2.

Q: Are you looking for a GRE type tunnel between switches?

A: No, use a Layer 2 switch tunneling feature.

Q: I assume you require remote span configured for R2 traffic. Is it okay to send both TX and RX traffic to Switch 2?

A: Read the question carefully because this information has been provided.

Q: Would you like me to configure UDLD aggressive mode on Switch 2 to transition the required port to error-disable mode if a trunk failure occurs?

A: No, you need to configure a feature that will place a nontrunk link into error-disable mode if all the trunks on Switch 2 fail.

Q: Can I just shut down ports 0/9 and 0/10 so that they can't communicate?

A: Nice try; look for a security feature to disable communication between these ports.

Section 2: IPv4 IGP Protocols

Section 2.1: EIGRP

Q: If I configure a summary-address on R1, this route overrides the VLAN100 and Loopback 0 routes from R1 as received on R3. Is this correct?

A: Yes, this is the expected behavior of summarization; you need to enable a feature that enables the more specific routes to be received on R3.

Q: I think I can achieve this with multiple summary routes but the question restricts this. Can I use a new EIGRP process instead?

A: No, use a feature that enables your specific routes to leak from the summary route.

Q: Is it acceptable to adjust the hold-time on the Frame-Relay interfaces to change the hello-interval?

A: Yes.

Q: Can I manipulate the delay associated to network 150.101.1.0/24 because this advertisement leaves R2 rather than by changing an interface delay on R2?

A: Yes.

Section 2.2: OSPF

Q: I am experiencing neighbor adjacency issues between R5 and Switch 1. Is this part of the question?

A: This is a byproduct of the question if you use a 3550 in your topology. Practice your troubleshooting skills to determine what issues could be causing this behavior.

Q: I've checked my configs between R5 and Switch 1 and they look good. Am I missing something from the initial configuration?

A: No, if your configuration is correct, you should debug your adjacencies to provide information on what could be causing an issue.

Q: I've found an MTU issue while debugging. Is it okay to change an interface MTU to fix this issue?

A: Yes.

Q: Is it acceptable to provide tunnels between R4 and R5 to join area 0?

A: No, this solution would involve a neighbor relationship being formed between the routers in Area 0.

Q: I'd normally use a virtual link to extend Area 0 into a transit area. Can I use this technique to stretch Area 0 between R4 and R5?

A: You can use virtual links in your solution; think about where the links need to be though, to ensure your topology operates correctly.

Section 2.3: RIPv2

Q: I've just checked the routing table of R3 to find the only RIPv2 route received from R2 is the route required in the question. Can I move on or have I missed something?

A: Read the question again; even if you have only a single RIPv2 route in your routing table, it doesn't mean it is the only RIPv2 route received by R3.

Q: I can see that I am of course still generating additional routes from R2 toward R3. Can I just block these with a distribute-list on R3?

A: Yes.

Q: Can I just use the passive-interface feature on the interfaces on R3 to make sure they are not advertised to R2?

A: No, this would stop RIPv2 advertisements from being sent out on these interfaces; it wouldn't stop the actual interface subnets from being advertised to R2.

Q: Can I create an offset-list on R3 marking the attached networks on R3 as unreachable so that they are not advertised to R2?

A: No; look for a simple solution that blocks routing advertisements from leaving an interface.

Q: So I'm okay to use the passive-interface feature on the Frame-Relay interface to stop advertising outbound but still receive the specific route from R2 inbound?

A: Yes.

Section 2.3: Redistribution

Q: I've followed the redistribution instructions, but I don't receive the RIPv2 route on R1 after redistribution.

A: You will have some underlying issues prior to receiving the route on R1; use your troubleshooting skills to determine the problem.

Q: I've noticed that due to the preconfigured Loopback interfaces on R1 and R3 both of these routers have the same EIGRP router-id. Can I manually change the router-id on one of the routers to see if this helps?

A: Yes.

Q: I've managed to get the RIPv2 route redistributed from R3 into EIGRP on R1, but the next hop is showing as R3. Can I policy-route on R1 so that the next hop for this route is directly via R2?

A: No, you need to have the routing table reflect the next hop of this route via R2 and not R3.

Q: Can I use the eigrp third-party next-hop feature to leave the next hop of the route unaltered from R2?

A: Yes.

Q: Can I modify the OSPF cost on the interface connecting R3 to the OSPF network to attempt to change the next hop for the subnet 120.100.63.0/24?

A: No, this would affect routes received on R3 from both R4 and R5 equally because R4 and R5 reside on the same subnet as R3.

Q: Can I use an offset-list or similar feature on R4 to penalize the route 120.100.63.0/24 as it advertised to R3?

A: No, you are permitted to configure only R3.

Q: Is it acceptable to use a route-map on R3 and match a route source to penalize the route to 120.100.63.0/24?

A: Yes.

Section 3: BGP

Q: If I can't use ebgp-multihop on my peering on R6, Switch 3, and Switch 4, will my peering fail because I am peering from my Loopback interfaces?

A: Yes, it will; you need to configure a feature that overrides this behavior.

Q: Can I try to use NAT to fix my peering?

A: No, use a specific BGP feature to disregard the TTL check.

Q: I'm experiencing peering issues between R1 and R3 and have BGP notifications displayed on the console. Is this expected behavior?

A: Yes, you had a similar issue within EIGRP; check your router-id.

Q: Do you want me to configure an ACL to limit BGP connections to purely inbound or outbound on TCP port 179?

A: No, an ACL would actually break the peering entirely. Use a BGP feature to force the peering to become directional.

Q: Can I just configure a filter on R4 to stop advertising specific routes to R3?

A: No, you must dynamically inform R4 to not advertise specific routes via R3.

Q: Can I use BGP ORF?

A: Yes.

Section 4: IPv6

Q: Would you like me to configure an additional IPv6 subnet on R4 to receive the 2007::/16 route?

A: No, investigate an alternative method to create this route from the preconfigured subnets you already have, ensuring that the route is received as illustrated in the question.

Q: Would you like me to redistribute routes into OSPFv3 as External Type 1 or Type 2?

A: The question provides you with sufficient information to determine the redistribution type to use.

Section 5: QoS

Q: Do the VC units use UDP Ports 3230 and 3235 or 3230 through 3235?

A: They use the range 3230 through 3235.

Q: Do you want me to trust the ports assigned to the VC units?

A: The VC devices are not marking the traffic, so there is a need to trust these ports.

Q: Would you like me to disable trunking, channeling, and spanning-tree checks on the ports assigned to the VC units?

A: Yes, but remember there is a single command that will disable all these features.

Q: If I use the bandwidth percent command on R2 in my 40-percent guaranteed reservation, is this sufficient to answer the question?

A: No, the question dictates that a priority queue be used.

Q: Would you like me to configure RTP compression within a frame-relay map-class?

A: No, you can achieve all the requirements within the same QoS policy-map.

Section 6: Multicast

Q: If I configure R1 and R2 for the same multicast groups, won't R3 and R4 see both routers as RPs for the same groups?

A: Yes, you will address this behavior in the following question.

Q: To have R1 and R2 as candidate RPs for different groups, can I just configure group-lists on R3?

A: If you were permitted to configure R1 and R2, group-lists would achieve the desired results, but you are permitted to configure only R3. Group-lists can assist in your solution on R3, but you need to find a method of assigning these specifically to R1 and R2.

Q: Do you want me to actually configure an IGMP join-group on R1 for 225.225.0.1 for the SNMP question?

A: No, this isn't required; traffic destined to this group will be sent to R1 regardless because it is the candidate RP for this group.

Section 7: Security

Q: Do you want me to configure an ACL to block SYN packets coming into VLAN63?

A: No, SYN packets should still enter into VLAN63. You need to configure a feature that monitors the SYN packets and closes down any half-opened connections.

Q: Can I use a reflexive ACL to drop SYN packets that are not correctly established by the servers?

A: No, there is a specific TCP feature used to protect servers from a flood of SYN packets that could cause a DoS attack.

Q: Can I just use a standard ACL on R1 on the frame-relay interface to permit sessions outbound and deny everything else inbound?

A: No, this would block return path traffic initiated by R1.

Q: Can I use a reflexive ACL to dynamically permit the return traffic with a time limit of 100 seconds?

A: Yes.

Q: I have configured SCP with the required SSH parameters, but I am not confident of my configuration; any suggestions?

A: If you have time, try to copy the IOS image from flash on R1 with RCP. If you are prompted for a password and gain access to the file, you have configured this feature correctly.

Practice Lab Debrief

The lab debrief section now analyzes each question showing you what was required and how to achieve the desired results. You should use this section to produce an overall score for the practice lab.

Section 1: LAN Switching and Frame-Relay (24 Points)

- Configure your switched network to use 802.1w Spanning Tree. Switch 1 should be the root bridge for VLANs 34,46,53,63,100, and 200, with Switch 2 being the secondary root bridge for all listed VLANs. (3 points)

802.1w is a rapid spanning tree; the switches will be in the default mode of standard PVST and require configuration to rapid-pvst mode. Switch 1 is required to be the root bridge and Switch 2 the secondary root bridge for VLANs 34, 46, 53, 63, 100, and 200. If you have configured this correctly, as shown in Example 2-1, you have earned 3 points. Example 2-1 also shows confirmation of the root bridge and which interfaces are used to reach the root bridge from the neighboring switches, VLAN 34 is used as an example but each VLAN would be identical in this configuration.

EXAMPLE 2-1 Sw1, Sw2, Sw3 and Sw4 Configuration and Verification

```
SW1(config)# spanning-tree mode rapid-pvst
SW1(config)# spanning-tree vlan 34,46,53,63,100,200 root primary
```

```
SW2 (config)# spanning-tree mode rapid-pvst
SW2 (config)# spanning-tree vlan 34,46,53,63,100,200 root secondary

SW3 (config)# spanning-tree mode rapid-pvst

SW4 (config)# spanning-tree mode rapid-pvst

SW1# show spanning-tree vlan 34 | include root
      This bridge is the root
SW1# show spanning-tree vlan 46 | include root
      This bridge is the root
SW1# show spanning-tree vlan 53 | include root
      This bridge is the root
SW1# show spanning-tree vlan 63 | include root
      This bridge is the root
SW1# show spanning-tree vlan 100 | include root
      This bridge is the root

SW1# show spanning-tree vlan 200 | include root
      This bridge is the root

SW2# show spanning-tree vlan 34 | include Root FWD
Fa0/23          Root FWD 19          128.25   P2p

SW3# show spanning-tree vlan 34 | include Root FWD
Fa0/19          Root FWD 19          128.21   P2p

SW4# show spanning-tree vlan 34 | include Root FWD
Fa0/21          Root FWD 19          128.23   P2p
```

- Switch 3 should use its interface directly connecting to Switch 2 (Fast Ethernet 0/21) for traffic directed toward even-numbered VLANs (34, 46, 100, 200) and the interface directly connecting to Switch 1 (Fast Ethernet 0/19) for odd-numbered VLANs (53, 63). (3 points)

This is a straightforward VLAN load-balancing question to ensure that trunk links are utilized efficiently and not logically disabled by spanning tree. Switch 3 uses the interface directly connecting to Switch 1 (Fast Ethernet 0/19) for all

VLANs as the lowest root cost path by default. To adjust this behavior, this interface must effectively be penalized for the even-numbered VLANs to ensure a more attractive path is via Switch 2 (Fast Ethernet 0/21). If you have configured this correctly, as shown in Example 2-2, you have scored 3 points.

EXAMPLE 2-2 Sw3 VLAN Load Balancing Configuration and Verification

```
SW3(config)# interface FastEthernet 0/19
SW3(config-if)# spanning-tree vlan 34,46,100,200 cost 100
```

```
SW3(config-if)# do show spanning-tree root
```

Vlan	Root ID	Root Cost	Hello Time	Max Age	Fwd Dly	Root Port
VLAN0001	32769 0013.806d.9400	19	2	20	15	Fa0/19
VLAN0034	24610 0013.806d.9400	38	2	20	15	
VLAN0046	24622 0013.806d.9400	38	2	20	15	
VLAN0053	24629 0013.806d.9400	19	2	20	15	Fa0/19
VLAN0063	24639 0013.806d.9400	19	2	20	15	Fa0/19
VLAN0100	24676 0013.806d.9400	38	2	20	15	Fa0/21
VLAN0200	24776 0013.806d.9400	38	2	20	15	Fa0/21

- Switch 4 should use its interface directly connecting to Switch 2 (Fast Ethernet0/19) for traffic destined toward even-numbered VLANs (34, 46, 100, 200) and the interface directly connected to Switch 1 (Fast Ethernet 0/21) for odd-numbered VLANs (53, 63). (3 points)

Following from the previous question, to ensure a balanced access topology for VLAN load balancing, Switch 4 uses the interface directly connecting to Switch 1 (Fast Ethernet 0/21) for all VLANs as the lowest root cost path by default, rendering the second trunk connecting to Switch 2 unused unless a failover condition occurs. As per the previous question, the directly connected interface to Switch 1 needs to be penalized for the even-numbered VLANs. If you have configured this correctly, as shown in Example 2-3, you have scored 3 points.

EXAMPLE 2-3 Sw4 VLAN Load Balancing Configuration and Verification

```
SW4(config)# interface FastEthernet 0/21
SW4(config-if)# spanning-tree vlan 34,46,100,200 cost 100
SW4(config-if)# do show spanning-tree root
```

```
Root Hello Max Fwd
```

Vlan	Root ID	Cost	Time	Age	Dly	Root Port
VLAN0001	32769 0013.806d.9400	19	2	20	15	Fa0/21
VLAN0034	24610 0013.806d.9400	38	2	20	15	Fa0/19
VLAN0046	24622 0013.806d.9400	38	2	20	15	Fa0/19
VLAN0053	24629 0013.806d.9400	19	2	20	15	Fa0/21
VLAN0063	24639 0013.806d.9400	19	2	20	15	Fa0/21
VLAN0100	24676 0013.806d.9400	38	2	20	15	Fa0/19
VLAN0200	24776 0013.806d.9400	38	2	20	15	Fa0/19

- Ensure that a cable fault between Switches 1 and 2 could not result in one-way traffic between the two switches, resulting in spanning-tree issues.(2 points)

UDLD detects unidirectional links on fiber-optic connections, in aggressive mode. UDLD also detects unidirectional links because of one-way traffic on twisted-pair links. By configuring the ports between Switch 1 and Switch 2 into aggressive mode, the switches become UDLD neighbors, can detect one-way links, and shut down the link if this condition arises to mitigate spanning-tree issues. If you have configured this correctly, as shown in Example 2-4, you have scored 2 points.

EXAMPLE 2-4 Sw1 and Sw2 UDLD Configuration and Verification

```

SW1(config)# interface FastEthernet 0/23
SW1(config-if)# udld port aggressive

SW2(config)# interface FastEthernet 0/23
SW2(config-if)# udld port aggressive

SW1# show udld FastEthernet 0/23

Interface Fa0/23
---
Port enable administrative configuration setting: Enabled / in aggressive mode
Port enable operational state: Enabled / in aggressive mode
Current bidirectional state: Bidirectional
Current operational state: Advertisement - Single neighbor detected
Message interval: 15
Time out interval: 5

Entry 1

```

```
---
Expiration time: 44
Cache Device index: 1
Current neighbor state: Bidirectional
Device ID: CAT0935N2GQ
Port ID: Fa0/23
Neighbor echo 1 device: CAT0911X17K
Neighbor echo 1 port: Fa0/23

Message interval: 15
Time out interval: 5
CDP Device name: SW2
```

- Configure Switch 1 and Switch 2 to allow connectivity of two further switches in the future to be connected to ports Fast Ethernet 0/18 on each switch. The new switches should be able to tunnel their own configured VLANs through a new VLAN (30) between Switch 1 and Switch 2. There is no requirement to configure a root bridge or VLAN load balancing for the new VLAN between Switch 1 and Switch 2. (4 points)

This is a service provider requirement whereby customers tunnel their own VLANs through the providers network; To mitigate any VLAN overlaps from other customers, a unique service provider VLAN is used to transport the customer VLANs. Example 2-5 shows VLAN 30 being used to transport VLANs over a dot1q-tunnel. Use the **show dot1q-tunnel** command to verify your tunnel configuration on your switches. If you have configured this correctly, as shown in Example 2-5, you have scored 4 points.

EXAMPLE 2-5 Sw1 and Sw2 Q in Q Configuration

```
SW1(config)# vlan 30
SW1(config-vlan)# exit
SW1(config)# interface FastEthernet 0/18
SW1(config-if)# switchport access vlan 30
SW1(config-if)# switchport mode dot1q-tunnel

SW2(config)# vlan 30
SW2(config-vlan)# exit
SW2(config)# interface FastEthernet 0/18
SW2(config-if)# switchport access vlan 30
SW2(config-if)# switchport mode dot1q-tunnel
```

- Configure your switched network to monitor the VLAN200 interface associated with R2 (Switch 2 FastEthernet 0/1) and send only traffic destined to R2 on this switch port across your network to Switch 3 port Fast Ethernet 0/17—use a new VLAN (20) to assist in this configuration. There is no requirement to configure a root bridge or VLAN load balancing for the new VLAN. (3 points)

This is a remote span question; the only complexity is based around the question statement of where you actually need to monitor—“traffic destined to R2.” As such, this means you need to configure the span parameters to only send the traffic transmitted out of the switch port toward R2, which is configured by the TX parameter. If this optional parameter is not configured, both transmit and receive traffic is monitored. Remote span requires a VLAN to propagate the span traffic between switches, which is why you need to configure VLAN 20 on both Switches 1 and 2. If you have configured this correctly, as shown in Example 2-6, you have scored 3 points.

EXAMPLE 2-6 Sw2 and Sw2 Remote Span Configuration and Verification

```
SW2(config)# vlan 20
SW2(config-vlan)# remote-span
SW2(config-vlan)# exit
SW2(config)# monitor session 1 source interface fastEthernet 0/1 tx
SW2(config)# monitor session 1 destination remote vlan 20
SW2(config)# do show monitor session 1
Session 1
-----
Type                : Remote Source Session
Source Ports       :
TX Only            : Fa0/1
Dest RSPAN VLAN   : 20

SW3(config)# vlan 20
SW3(config-vlan)# exit
SW3(config)# monitor session 1 source remote vlan 20
SW3(config)# monitor session 1 destination interface fast 0/17
SW3(config)# do show monitor session 1
Session 1
-----
Type                : Remote Destination Session
Source RSPAN VLAN  : 20
Destination Ports  : Fa0/17
```

```
Encapsulation : Native
Ingress       : Disabled
```

- Configure the interface on Switch 2, which connects to R5 VLAN53 (Fast Ethernet 0/5) in such a way that if all the trunks on Switch 2 connecting to Switch 1, Switch 3, and Switch 4 should fail, this Ethernet port transitions into error-disable state. (3 points)

The question requires link-state tracking to be configured. This feature provides redundancy in the network when used with server NIC adapter teaming. If a link is lost on the primary interface, connectivity is transparently switched to the secondary interface. Ports connected to servers are configured as downstream ports, and ports connected to other switches are configured as upstream ports. If the upstream trunk ports on Switch 2 fail, link-state tracking automatically puts the downstream port connected to R5 into error-disable state. Example 2-7 shows the associated configuration and testing by shutting down the trunk ports on Switch 2, which connects to Switch 1, Switch 3, and Switch 4, which forces FastEthernet downstream port into error-disable state. If you have configured this correctly, as shown in Example 2-7, you have scored 3 points.

EXAMPLE 2-7 Sw2 Link-State Tracking Configuration and Verification

```
SW2(config)# link state track 1
SW2(config)# interface fast0/5
SW2(config-if)# link state group 1 downstream
SW2(config-if)# interface FastEthernet 0/19
SW2(config-if)# link state group 1 upstream
SW2(config-if)# interface FastEthernet 0/21
SW2(config-if)# link state group 1 upstream
SW2(config-if)# interface FastEthernet 0/23
SW2(config-if)# link state group 1 upstream

SW2# show interface FastEthernet 0/5 | include connected
FastEthernet0/5 is up, line protocol is up (connected)

SW2(config-if)# int fast 0/19
SW2(config-if)# shut
SW2(config-if)# int fast 0/21
SW2(config-if)# shut
SW2(config-if)# int fast 0/23
SW2(config-if)# shut
```

```
SW2# show interface FastEthernet 0/5 | include err-disabled
FastEthernet0/5 is down, line protocol is down (err-disabled)
```

- Configure interfaces Fast Ethernet 0/9 and 0/10 on Switch 1 so that even if they are configured to belong to the same VLAN they cannot forward unicast, broadcast, or multicast traffic to one another. Do not use any form of ACL or configure the ports to belong to a PVLAN. (1 point)

You are required to configure the interfaces with the command **switchport protected** to ensure that no traffic is forwarded between these ports. Traffic is forwarded as normal between a protected and nonprotected port. If you have configured this correctly, you have scored 1 point.

- Your initial Frame-Relay configuration has been supplied for the R1-R2-R3 connectivity. Configure Frame-Relay as per Figure 2-6 to ensure each device is reachable over the Frame-Relay network. Use only the indicated DLCIs and ensure that a proprietary method of reducing the payload over the Frame-Relay network is enabled on a per packet basis. (2 points)

The initial Frame-Relay configuration has been supplied for you; all you need to add is additional maps on R1 and R2 spokes to enable them to communicate with each other by directing traffic to the hub router (R3) as the initial configuration uses no inverse ARP. To reduce the payload, you are required to enable payload-compression packet-by-packet within the map statements. If you have configured this correctly, as shown in Example 2-8, you have scored 2 points.

EXAMPLE 2-8 R1 and R2 Additional Frame-Relay Configuration and Testing

```
R1(config)# interface Serial0/0/0
R1(config-if)# frame-relay map ip 120.100.123.2 103 broadcast payload-compression packet-by-packet

R2(config)# interface Serial0/0
R2(config-if)# frame-relay map ip 120.100.123.1 203 broadcast payload-compression packet-by-packet

R3(config)# interface Serial0/0
R3(config-if)# frame-relay map ip 120.100.123.1 301 broadcast payload-compression packet-by-packet
R3(config-if)# frame-relay map ip 120.100.123.2 302 broadcast payload-compression packet-by-packet

R1# ping 120.100.123.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 120.100.123.2, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 8/8/8 ms
```

Section 2: IPv4 IGP Protocols (28 Points)

Section 2.1: EIGRP

- Configure EIGRP per Figure 2-7 using an AS of 1. Each EIGRP router should have its Loopback 0 interface configured and advertised within EIGRP. (2 points)

Use vanilla EIGRP configuration in preparation for the following questions, the only complexity is spotting the split-horizon issue with R3 over the physical frame-relay network. By default, R3 will not advertise the routes learned on its Serial interface from R1 back out to R2 and vice versa because they all share the same interface. By disabling split-horizon for EIGRP on R3, the routes are permitted to propagate. If you have configured this correctly, as shown in Example 2-9, you have scored 2 points.

EXAMPLE 2-9 EIGRP Configuration and Verification

```
R1(config)# router eigrp 1
R1(config-router)# no auto-summary
R1(config-router)# net 120.100.1.0 0.0.0.255
R1(config-router)# net 120.100.123.0 0.0.0.255
R1(config-router)# net 120.100.100.0 0.0.0.255

R2(config)# router eigrp 1
R2(config-router)# no auto-summary
R2(config-router)# network 120.100.2.0 0.0.0.255
R2(config-router)# network 120.100.123.0 0.0.0.255
R2(config-router)# network 120.100.200.0 0.0.0.255

R3(config-if)# router eigrp 1
R3(config-router)# no auto-summary
R3(config-router)# network 120.100.3.0 0.0.0.255
R3(config-router)# network 120.100.123.0 0.0.0.255
R3(config-router)# network 120.100.34.0 0.0.0.255

R4(config-router)# router eigrp 1
R4(config-router)# no auto-summary
R4(config-router)# network 120.100.4.0 0.0.0.255
R4(config-router)# network 120.100.34.0 0.0.0.255
```

```
R5(config)# router eigrp 1
R5(config-router)# no auto-summary
R5(config-router)# network 120.100.5.0 0.0.0.255
R5(config-router)# network 120.100.34.0 0.0.0.255

R1# show ip route eigrp
120.0.0.0/8 is variably subnetted, 7 subnets, 1 mask
D       120.100.4.0/24 [90/2300416] via 120.100.123.3, 00:14:51, Serial0/0/0
D       120.100.5.0/24 [90/2300416] via 120.100.123.3, 00:01:32, Serial0/0/0
D       120.100.3.0/24 [90/2297856] via 120.100.123.3, 00:42:12, Serial0/0/0
D       120.100.34.0/24 [90/2172416] via 120.100.123.3, 00:41:54, Serial0/0/0

R3# show ip route eigrp
120.0.0.0/8 is variably subnetted, 9 subnets, 1 mask
D       120.100.4.0/24
        [90/156160] via 120.100.34.4, 00:19:14, GigabitEthernet0/0
D       120.100.5.0/24
        [90/156160] via 120.100.34.5, 00:05:55, GigabitEthernet0/0
D       120.100.1.0/24 [90/2297856] via 120.100.123.1, 00:46:35, Serial0/0/0
D       120.100.2.0/24 [90/2297856] via 120.100.123.2, 00:46:35, Serial0/0/0
D       120.100.100.0/24 [90/2172416] via 120.100.123.1, 00:46:35, Serial0/0/0
D       120.100.200.0/24 [90/2172416] via 120.100.123.2, 00:46:35, Serial0/0/0

R2# show ip route eigrp
120.0.0.0/8 is variably subnetted, 8 subnets, 1 mask
D       120.100.4.0/24 [90/2300416] via 120.100.123.3, 00:19:55, Serial0/0
D       120.100.5.0/24 [90/2300416] via 120.100.123.3, 00:06:36, Serial0/0
D       120.100.3.0/24 [90/2297856] via 120.100.123.3, 00:47:16, Serial0/0
D       120.100.34.0/24 [90/2172416] via 120.100.123.3, 00:46:58, Serial0/0

R3(config)# interface Serial0/0/0
R3(config-if)# no ip split-horizon eigrp 1

R1# show ip route eigrp
120.0.0.0/8 is variably subnetted, 10 subnets, 1 mask
D       120.100.4.0/24 [90/2300416] via 120.100.123.3, 00:14:51, Serial0/0/0
D       120.100.5.0/24 [90/2300416] via 120.100.123.3, 00:01:32, Serial0/0/0
D       120.100.2.0/24 [90/2809856] via 120.100.123.3, 00:38:32, Serial0/0/0
```

```

D      120.100.3.0/24 [90/2297856] via 120.100.123.3, 00:42:12, Serial0/0/0
D      120.100.34.0/24 [90/2172416] via 120.100.123.3, 00:41:54, Serial0/0/0
D      120.100.200.0/24 [90/2684416] via 120.100.123.3, 00:38:32, Serial0/0/0
R1#

R2# show ip route eigrp
      120.0.0.0/8 is variably subnetted, 10 subnets, 1 mask
D      120.100.4.0/24 [90/2300416] via 120.100.123.3, 00:24:43, Serial0/0
D      120.100.5.0/24 [90/2300416] via 120.100.123.3, 00:11:24, Serial0/0
D      120.100.1.0/24 [90/2809856] via 120.100.123.3, 00:48:24, Serial0/0
D      120.100.3.0/24 [90/2297856] via 120.100.123.3, 00:52:04, Serial0/0
D      120.100.34.0/24 [90/2172416] via 120.100.123.3, 00:51:46, Serial0/0
D      120.100.100.0/24 [90/2684416] via 120.100.123.3, 00:48:24, Serial0/0

```

- Configure R1 to advertise a summary route of 120.100.0.0/16 outbound on its serial interface. R3 should see the original VLAN100 and Loopback 0 individual routes in addition to the summary route. You can only use one summary route in your configuration. (3 points)

Summarization will by default block all longer prefixes covered by the supernet configured on an interface; as such, the VLAN 100 and Loopback 0 route from R1 would not be seen by R3. Allowing specific routes to be advertised with summary routes can be a valid requirement. One method used to achieve this is by configuring multiple summary routes, but the question does not permit this approach. To facilitate the specific routes with the summary, a leak-map should be configured to match the VLAN 100 and Loopback 0 interfaces on R1. The leak-map, which is configured per a normal route-map, is then applied to the standard summary route statement on R1. If you have configured this correctly, as shown in Example 2-10, you have scored 3 points.

EXAMPLE 2-10 R1 Leak Map Configuration and Verification

```

R1(config)# route-map LEAK-VLAN-100-LOOP0 permit 10
R1(config-route-map)# match ip address 1
R1(config-route-map)# exit
R1(config)# access-list 1 permit 120.100.100.0
R1(config)# access-list 1 permit 120.100.1.0
R1(config)# interface Serial0/0/0
R1(config-if)# ip summary-address eigrp 1 120.100.0.0 255.255.0.0 leak-map LEAK-VLAN-100-LOOP0

R3# show ip route eigrp
      120.0.0.0/8 is variably subnetted, 10 subnets, 2 masks
D      120.100.4.0/24

```

```

          [90/156160] via 120.100.34.4, 00:19:14, GigabitEthernet0/0
D      120.100.5.0/24
          [90/156160] via 120.100.34.5, 00:05:55, GigabitEthernet0/0
D      120.100.0.0/16 [90/2172416] via 120.100.123.1, 00:34:39, Serial0/0/0
D      120.100.1.0/24 [90/2297856] via 120.100.123.1, 00:46:35, Serial0/0/0
D      120.100.2.0/24 [90/2297856] via 120.100.123.2, 00:46:35, Serial0/0/0
D      120.100.100.0/24 [90/2172416] via 120.100.123.1, 00:46:35, Serial0/0/0
D      120.100.200.0/24 [90/2172416] via 120.100.123.2, 00:46:35, Serial0/0/0

```

- Ensure the length of time that EIGRP considers neighbors to be valid without receiving a hello packet on the Frame-Relay network between R1, R2, and R3 is 200 seconds; do not change the hello-interval parameter. (2 points)

EIGRP considers neighbors to be valid up to three times the hello interval, the Frame-Relay network is considered a slow speed link, and hello packets will be sent every 60 seconds. Usually you could tune the hold time by manipulating the hello intervals on an interface, but this question ensures you can achieve the desired result only by manually changing the hold-time to 200 under the Frame-Relay interface of Routers R1, R2, and R3. Example 2-11 shows the required configuration and verification of hold time by displaying the neighbors' statistics as seen by R3. If you have configured this correctly, as shown in Example 2-11, you have scored 2 points.

EXAMPLE 2-11 EIGRP Hold Timer Configuration and Verification

```

R1(config)# interface Serial0/0/0
R1(config-if)# ip hold-time eigrp 1 200
R1(config-if)

```

Enter configuration commands, one per line. End with CNTL/Z.

```

R2(config)# interface Serial0/0
R2(config-if)# ip hold-time eigrp 1 200
R2(config-if)

```

```

R3(config)# interface Serial0/0/0
R3(config-if)# ip hold-time eigrp 1 200
R3(config-if)# do sh ip eigrp neighbors

```

```

IP-EIGRP neighbors for process 1

```

H	Address	Interface	Hold (sec)	Uptime	SRTT (ms)	RTO	Q Cnt	Seq Num
3	120.100.123.1	Se0/0/0	198	00:00:57	3	200	0	25
2	120.100.123.2	Se0/0/0	199	00:01:00	3	200	0	18

1	120.100.34.5	Gi0/0	12 00:23:32	1	200	0	21
0	120.100.34.4	Gi0/0	12 00:23:35	35	210	0	22

- Configure new Loopback interfaces on R1 and R2 using a Loopback interface 2 with an identical IP address of 150.101.1.1/24 on both routers; advertise this network into EIGRP on each router. Ensure that R3 prefers the route from R2 by manipulating the delay associated with this route. Do not manually adjust the delay associated with the interface by use of the `delay` command, and you are permitted to configure only R2 to influence the delay. (3 points)

R3 will receive identical routes from both R1 and R2 for network 150.101.1.0/24; as such, both routes will be stored in the topology and routing table. R2 could influence the metric calculated by R3 by manipulating the delay of the new Loopback interface or of the serial Frame-Relay interface connecting directly to R3, but this is not permitted. As configuration is required solely on R2, the only method available is to create an offset-list, which enables you to match specific routes and append further delay to them as they are advertised on R2 toward R3. If the offset-list is not applied to the Frame-Relay interface, it would affect the whole process and not just advertisements toward R3. Example 2-12 shows the configuration required to advertise the new routes and the routes as they are received on R3. Initial delay is shown to be 25,000 μ S. Post configuration of the offset-list on R2, the delay is seen to increase to 25,003 μ S for the route received from R2; as such the route installed into the routing table of R3 is then the original advertised from R1 with the more appealing value of 25,000 μ S. If you have configured this correctly, as shown in Example 2-12, you have scored 3 points.

EXAMPLE 2-12 EIGRP Configuration and Verification

```
R1(config)# interface Loopback2
R1(config-if)# ip address 150.101.1.1 255.255.255.0
R1(config-if)# router eigrp 1
R1(config-router)# net 150.101.1.0 0.0.0.255

R2(config)# interface Loopback2
R2(config-if)# ip address 150.101.1.1 255.255.255.0
R2(config-if)# router eigrp 1
R2(config-router)# net 150.101.1.0 0.0.0.255

R3# show ip route 150.101.1.0
Routing entry for 150.101.1.0/24
  Known via "eigrp 1", distance 90, metric 2297856, type internal
  Redistributing via eigrp 1
```

```
Last update from 120.100.123.2 on Serial0/0/0, 00:02:51 ago
Routing Descriptor Blocks:
 120.100.123.2, from 120.100.123.2, 00:02:51 ago, via Serial0/0/0
   Route metric is 2297856, traffic share count is 1
   Total delay is 25000 microseconds, minimum bandwidth is 1544 Kbit
   Reliability 255/255, minimum MTU 1500 bytes
   Loading 1/255, Hops 1
 * 120.100.123.1, from 120.100.123.1, 00:02:51 ago, via Serial0/0/0
   Route metric is 2297856, traffic share count is 1
   Total delay is 25000 microseconds, minimum bandwidth is 1544 Kbit
   Reliability 255/255, minimum MTU 1500 bytes
   Loading 1/255, Hops 1
```

R3# **show ip eigrp topology 150.101.1.0 255.255.255.0**

```
IP-EIGRP (AS 1): Topology entry for 150.101.1.0/24
  State is Passive, Query origin flag is 1, 2 Successor(s), FD is 2297856
  Routing Descriptor Blocks:
 120.100.123.2 (Serial0/0/0), from 120.100.123.2, Send flag is 0x0
   Composite metric is (2297856/128256), Route is Internal
   Vector metric:
     Minimum bandwidth is 1544 Kbit
     Total delay is 25000 microseconds
     Reliability is 255/255
     Load is 1/255
     Minimum MTU is 1500
     Hop count is 1
 120.100.123.1 (Serial0/0/0), from 120.100.123.1, Send flag is 0x0
   Composite metric is (2297856/128256), Route is Internal
   Vector metric:
     Minimum bandwidth is 1544 Kbit
     Total delay is 25000 microseconds
     Reliability is 255/255
     Load is 1/255
     Minimum MTU is 1500
     Hop count is 1
```

R2(config-router)# **do show interface Serial0/0**

```
Serial0/0 is up, line protocol is up
  Hardware is GT96K Serial
  Internet address is 120.100.123.2/24
```

```
MTU 1500 bytes, BW 1544 Kbit, DLY 20000 usec,  
reliability 255/255, txload 1/255, rxload 1/255
```

```
R2(config)# access-list 1 permit 150.101.1.0
```

```
R2(config)# router eigrp 1
```

```
R2(config-router)# offset-list 1 out 100 Serial0/0
```

```
R3# show ip route 150.101.1.0
```

```
Routing entry for 150.101.1.0/24
```

```
Known via "eigrp 1", distance 90, metric 2297856, type internal
```

```
Redistributing via eigrp 1
```

```
Last update from 120.100.123.1 on Serial0/0/0, 00:00:18 ago
```

```
Routing Descriptor Blocks:
```

```
* 120.100.123.1, from 120.100.123.1, 00:00:18 ago, via Serial0/0/0
```

```
Route metric is 2297856, traffic share count is 1
```

```
Total delay is 25000 microseconds, minimum bandwidth is 1544 Kbit
```

```
Reliability 255/255, minimum MTU 1500 bytes
```

```
Loading 1/255, Hops 1
```

```
R3# show ip eigrp topology 150.101.1.0 255.255.255.0
```

```
IP-EIGRP (AS 1): Topology entry for 150.101.1.0/24
```

```
State is Passive, Query origin flag is 1, 1 Successor(s), FD is 2297856
```

```
Routing Descriptor Blocks:
```

```
120.100.123.1 (Serial0/0/0), from 120.100.123.1, Send flag is 0x0
```

```
Composite metric is (2297856/128256), Route is Internal
```

```
Vector metric:
```

```
Minimum bandwidth is 1544 Kbit
```

```
Total delay is 25000 microseconds
```

```
Reliability is 255/255
```

```
Load is 1/255
```

```
Minimum MTU is 1500
```

```
Hop count is 1
```

```
120.100.123.2 (Serial0/0/0), from 120.100.123.2, Send flag is 0x0
```

```
Composite metric is (2297956/128356), Route is Internal
```

```
Vector metric:
```

```
Minimum bandwidth is 1544 Kbit
```

```
Total delay is 25003 microseconds
```

```
Reliability is 255/255
```

```
Load is 1/255
```

```
Minimum MTU is 1500
Hop count is 1
```

Section 2.2: OSPF

- Configure OSPF per Figure 2-8 using a process ID of 1; all OSPF configuration where possible should not be configured under the process ID. Each OSPF router should also have its Loopback 0 interface configured and advertised within OSPF as follows: (2 points)

R4 Loopback 0 – Area 0

R5 Loopback 0 – Area 0

R6 Loopback 0 – Area 1

Sw1 Loopback 0 – Area 2

Sw2 Loopback 0 – Area 1

Sw3 Loopback 0 – Area 2

Sw4 Loopback 0 – Area

As per Lab 1, the question directs you to configure OSPF directly under the interfaces of the routers; the switches still require configuration under the OSPF process running this version of IOS. Did you notice that Area 0 is partitioned? If you have configured this correctly, as shown in Example 2-13, you have scored 2 points. Consider using the **show ip ospf interface** command to verify your configuration.

EXAMPLE 2-13 Initial OSPF Configuration

```
R4(config)# interface Loopback 0
R4(config-if)# ip ospf 1 area 0
R4(config-if)# exit
R4(config)# interface GigabitEthernet 0/1
R4(config-if)# ip ospf 1 area 1

R5(config)# interface Loopback 0
R5(config-if)# ip ospf 1 area 0
R5(config-if)# exit
```

```
R5(config)# interface GigabitEthernet 0/1
R5(config-if)# ip ospf 1 area 2

R6(config)# interface Loopback 0
R6(config-if)# ip ospf 1 area 1
R6(config-if)# interface GigabitEthernet 0/0
R6(config-if)# ip ospf 1 area 1
R6(config-if)# interface GigabitEthernet 0/1
R6(config-if)# ip ospf 1 area 3

SW1(config)# ip routing
SW1(config)# router ospf 1
SW1(config-router)# network 120.100.7.1 0.0.0.0 area 2
SW1(config-router)# network 120.100.53.1 0.0.0.0 area 2

SW2(config)# ip routing
SW2(config-if)# router ospf 1
SW2(config-router)# net 120.100.8.1 0.0.0.0 area 1
SW2(config-router)# net 120.100.46.2 0.0.0.0 area 1

SW3(config)# ip routing
SW3(config)# router ospf 1
SW3(config-router)# network 120.100.53.3 0.0.0.0 area 2
SW3(config-router)# network 120.100.63.3 0.0.0.0 area 3
SW3(config-router)# network 120.100.9.1 0.0.0.0 area 2

SW4(config)# ip routing
SW4(config)# router ospf 1
SW4(config-router)# network 120.100.10.1 0.0.0.0 area 3
SW4(config-router)# network 120.100.63.4 0.0.0.0 area 3
```

If you are using a 3550 as one of your switches, you will experience neighbor relationship problems running OSPF to your routers or 3560s. This is because the default MTU value is 1504 on the 3550 VLAN interface and 1500 on the routers and 3560s. Example 2-14 shows the adjacency issues with Switch 1 (3550 in this scenario) on R5; by debugging OSPF adjacency it can be seen that Switch 1 has a larger default MTU, which will ensure the neighbor adjacency is only ever partial. The example also shows the Switch 3 (3560) default MTU value on the same VLAN 53 and the MTU modification required on Switch 1. No extra points if you needed to configure this workaround. If you didn't spot this,

you would lose points in this section because of not having full neighbor adjacencies on Switch 1. This type of issue shows just how important it is to constantly validate your configurations rather than simply expecting everything to work.

EXAMPLE 2-14 R5-Sw1 OSPF Neighbor Issues

```
R5# show ip ospf neighbor

Neighbor ID      Pri   State           Dead Time   Address      Interface
120.100.7.1      1     EXSTART/DROTHER 00:00:35    120.100.53.1 GigabitEtherne
t0/1
120.100.9.1      1     FULL/DR          00:00:38    120.100.53.3 GigabitEtherne
t0/1
R5# debug ip ospf adjacency
*May  8 20:38:41.059: OSPF: Nbr 120.100.7.1 has larger interface MTU
R5#

R5# show interface GigabitEthernet0/0 | begin MTU
MTU 1500 bytes, BW 100000 Kbit, DLY 100 usec,

SW1# show interface vlan 53 | begin MTU
MTU 1504 bytes, BW 1000000 Kbit, DLY 10 usec

SW3# show interface vlan 53 | beg MTU
MTU 1500 bytes, BW 1000000 Kbit, DLY 10 usec,

SW1(config-if)# int vlan 53
SW1(config-if)# ip mtu 1500

R5# show ip ospf neighbor

Neighbor ID      Pri   State           Dead Time   Address      Interface
120.100.7.1      1     FULL/DROTHER    00:00:34    120.100.53.1 GigabitEthernet0/1
120.100.9.1      1     FULL/DR          00:00:37    120.100.53.3 GigabitEthernet0/1
```

- Area 0 is partitioned between R4 and R5—ensure your network can accommodate this issue. You are not permitted to form any Area 0 neighbor relationship directly between R4 and R5 to join Area 0. (4 points)

A fundamental rule of OSPF is not to design your network with a partitioned backbone Area 0 or partition if of a failure condition occurs. A virtual-link between R4 and R5 would not work here because you would need to transit multiple OSPF areas. A tunnel between the two routers is also not permitted because this would form a direct neighbor relationship. You are required to configure a virtual-link between R5 and Switch 3 to propagate Area 3 routes and similarly between R4 and R6. By then creating an additional virtual-link between R6 and Switch 3, the two effective halves of the network have been joined at an Area 0 level. Remember to configure all virtual-links to the router ID of the remote router as opposed to the physical IP address on the corresponding interface. Example 2-15 shows the required configuration to create virtual-links between R5-SW3, R4-R6, and R6-SW3. The resulting routing table verification on Switch 4 shows all networks are being learned correctly post configuration. If you have configured this correctly, as shown in Example 2-15, you have scored 4 points.

EXAMPLE 2-15 OSPF Virtual-Link Configuration and Routing Table Verification

```
R5(config)# router ospf 1
R5(config-router)# area 2 virtual-link 120.100.9.1

SW3(config-router)# router ospf 1
SW3(config-router)# area 2 virtual-link 120.100.5.1

R4(config)# router ospf 1
R4(config-router)# area 1 virtual-link 120.100.6.1

R6(config-if)# router ospf 1
R6(config-router)# area 1 virtual-link 120.100.4.1
R6(config-router)# area 3 virtual-link 120.100.9.1

SW3(config-if)# router ospf 1
SW3(config-router)# area 3 virtual-link 120.100.6.1

SW4# sh ip route ospf
120.0.0.0/8 is variably subnetted, 10 subnets, 2 masks
O IA 120.100.9.1/32 [110/2] via 120.100.63.3, 00:00:54, Vlan63
O IA 120.100.8.1/32 [110/3] via 120.100.63.6, 00:00:54, Vlan63
O IA 120.100.5.1/32 [110/3] via 120.100.63.3, 00:00:54, Vlan63
O IA 120.100.4.1/32 [110/3] via 120.100.63.6, 00:00:54, Vlan63
O IA 120.100.7.1/32 [110/3] via 120.100.63.3, 00:00:54, Vlan63
O IA 120.100.6.1/32 [110/2] via 120.100.63.6, 00:00:54, Vlan63
```

```
o IA    120.100.53.0/24 [110/2] via 120.100.63.3, 00:00:54, Vlan63
o IA    120.100.46.0/24 [110/2] via 120.100.63.6, 00:00:55, Vlan63
```

Section 2.2: RIPv2

- Configure RIPv2 between R2 and R3, configure a new Loopback interface on R2 (Loopback 3) with an IP address of 150.101.2.1/24, and advertise this and only this network to R3 from R2. (2 points)

Although RIPv2 is capable of VLSM, it is nevertheless based on a classful protocol that will by default advertise all the connected interfaces of both R2 and R3 when the classful network command is used to activate the routing process. To restrict advertisement to solely the new Loopback interface from R2, a basic distribute-list is required. This should be applied either on the entire process or just on the Frame-Relay interface connecting to R3. It should permit only the new Loopback subnet of 150.101.2.0/24. If you're low on time, you may check the routing table of R3 to find that the only RIPv2 route received is that of the new Loopback 3 interface on R2. This is because the VLAN 200 and Loopback 0 interfaces of R2 already being learned via EIGRP, which of course has a lower admin distance and will therefore not be listed as RIPv2 routes within the routing table. Example 2-16 shows the basic RIPv2 configuration on R2 and R3 with debug of RIPv2 updates on R2 to illustrate which routes are being advertised to R3. The required distribute-list configuration is also shown. If you have configured this correctly, as shown in Example 2-16, you have scored 2 points.

EXAMPLE 2-16 R2 and R3 RIPv2 Configuration and Verification

```
R2(config)# interface Loopback3
R2(config-if)# ip add 150.101.2.1 255.255.255.0
R2(config-if)# router rip
R2(config-router)# version 2
R2(config-router)# no auto-summary
R2(config-router)# network 150.101.0.0
R2(config-router)# network 120.0.0.0

R3(config)# router rip
R3(config-router)# version 2
R3(config-router)# no auto-summary
R3(config-router)# network 120.0.0.0
R3(config-router)# do show ip route rip
      150.101.0.0/24 is subnetted, 2 subnets
R       150.101.2.0 [120/1] via 120.100.123.2, 00:00:05, Serial0/0/0

R2# sh ip route rip
```

```

R2# debug ip rip
*May  8 05:00:22.147:  RIP: sending v2 update to 224.0.0.9 via Serial10/0 (120.100
.123.2)
*May  8 05:00:22.147:  RIP: build update entries
*May  8 05:00:22.147:  120.100.2.0/24 via 0.0.0.0, metric 1, tag 0
*May  8 05:00:22.147:  120.100.123.0/24 via 0.0.0.0, metric 1, tag 0
*May  8 05:00:22.147:  120.100.200.0/24 via 0.0.0.0, metric 1, tag 0
*May  8 05:00:22.147:  150.101.1.0/24 via 0.0.0.0, metric 1, tag 0
*May  8 05:00:22.147:  150.101.2.0/24 via 0.0.0.0, metric 1, tag 0

R2(config)# router rip
R2(config-router)# distribute-list 2 out Serial10/0
R2(config-router)# exit
R2(config)# access-list 2 permit 150.101.2.0
R2(config)# exit
*May  8 05:02:40.271:  RIP: sending v2 update to 224.0.0.9 via Serial10/0 (120.100
.123.2)
*May  8 05:02:40.271:  RIP: build update entries
*May  8 05:02:40.271:  150.101.2.0/24 via 0.0.0.0, metric 1, tag 0
R2#

R3# show ip route rip
      150.101.0.0/24 is subnetted, 2 subnets
R       150.101.2.0 [120/1] via 120.100.123.2, 00:00:02, Serial10/0/0

```

- R3 should not advertise any connected interfaces into RIPv2. Do not filter routing advertisements to achieve this behavior. (2 points)

Because you are not permitted to filter routes as per the previous question, you simply configure the Frame-Relay interfaces to be passive on R3. This allows routing updates to be received inbound but stops routing advertisements outbound. Example 2-17 shows the RIPv2 routes advertised originally from R3 being received by R2 with the required configuration for R3; if you have configured this correctly, you have scored 2 points.

EXAMPLE 2-17 R3 RIPv2 Configuration and Verification

```

R2# debug ip rip
*May  8 05:05:10.031:  RIP: received v2 update from 120.100.123.3 on Serial10/0
*May  8 05:05:10.031:  120.100.3.0/24 via 0.0.0.0 in 1 hops
*May  8 05:05:10.031:  120.100.34.0/24 via 0.0.0.0 in 1 hops

```

```
*May  8 05:05:10.031: 120.100.123.0/24 via 0.0.0.0 in 1 hops

R3(config)#  router rip
R3(config-router)#  passive-interface Serial0/0/0
```

Section 2.3: Redistribution

- Perform a one-way redistribution of RIPv2 into EIGRP on R3 using the following default metric: 1544 20000 255 1 1500. Ensure that R1 shows a next hop for the RIPv2 advertised route of 150.101.2.0/24 of R2. Perform configuration only on R3 for this task. (3 points)

A simple redistribution question, on inspection you'd believe the only complexity would be that of modifying the next hop attribute for R1, which would by default show as R3 for the RIPv2 route advertised by R2. In fact, you would find that the RIPv2 route would not be seen on R1 post redistribution from R3. This is due to an inherent safety mechanism within EIGRP that will cause redistribution issues with routers that have duplicate EIGRP router IDs. Pre-lab configuration ensured that both R1 and R2 have the same Loopback 255 IP address, which will force the router ID to be identical. Example 2-18 shows the redistribution configuration on R3. The RIPv2 route of 150.101.2.0/24 is received on R3 but is absent on R1. Inspection of the EIGRP topology table for the route on R3 shows that it is being advertised into EIGRP and that the router ID of R3 is 200.200.200.200; similarly, the router ID of R1 is also 200.200.200.200. By changing the router ID of R3 to that of its Loopback 0 interface (120.100.3.1), the route is then accepted by R1, but of course a next hop is shown as R3, even though R2 resides on the same IP subnet as R1 and R2 and is the originating router. The EIGRP third-party next-hop feature can be used to modify the next-hop attribute with a router redistributing another routing protocol into EIGRP in a similar manner to that of BGP. If you have configured this correctly, as shown in Example 2-18, you have scored 3 points.

EXAMPLE 2-18 R3 RIPv2 Redistribution Configuration and Verification

```
R3(config)#  router eigrp 1
R3(config-router)#  redistribute rip
R3(config-router)#  default-metric 1544 20000 255 1 1500

R3#  show ip route rip
      150.101.0.0/24 is subnetted, 2 subnets
R       150.101.2.0 [120/1] via 120.100.123.2, 00:00:05, Serial0/0/0

R1#  show ip route 150.101.2.0
% Subnet not in table
```

```
R3# show ip eigrp topology 150.101.2.0 255.255.255.0
IP-EIGRP (AS 1): Topology entry for 150.101.2.0/24
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 6777856
  Routing Descriptor Blocks:
    120.100.123.2, from Redistributed, Send flag is 0x0
      Composite metric is (6777856/0), Route is External
      Vector metric:
        Minimum bandwidth is 1544 Kbit
        Total delay is 200000 microseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 0
      External data:
        Originating router is 200.200.200.200 (this system)
        AS number of route is 0
        External protocol is RIP, external metric is 1
        Administrator tag is 0 (0x00000000)
```

```
R3# show ip eigrp topology |include ID
IP-EIGRP Topology Table for AS(1)/ID(200.200.200.200)
R3#
```

```
R1# show ip eigrp topology |include ID
```

```
R1#
```

```
IP-EIGRP Topology Table for AS(1)/ID(200.200.200.200)
```

```
R3(config)# router eigrp 1
R3(config-router)# eigrp router-id 120.100.3.1
```

```
R3# show ip eigrp topology |include ID
IP-EIGRP Topology Table for AS(1)/ID(120.100.3.1)
```

```
R3# show ip eigrp topology 150.101.2.0 255.255.255.0
IP-EIGRP (AS 1): Topology entry for 150.101.2.0/24
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 6777856
  Routing Descriptor Blocks:
    120.100.123.2, from Redistributed, Send flag is 0x0
      Composite metric is (6777856/0), Route is External
```

```
Vector metric:
  Minimum bandwidth is 1544 Kbit
  Total delay is 200000 microseconds
  Reliability is 255/255
  Load is 1/255
  Minimum MTU is 1500
  Hop count is 0
External data:
  Originating router is 120.100.3.1 (this system)
  AS number of route is 0
  External protocol is RIP, external metric is 1
  Administrator tag is 0 (0x00000000)
```

R1# **show ip route 150.101.2.0**

```
Routing entry for 150.101.2.0/24
  Known via "eigrp 1", distance 170, metric 7289856, type external
  Redistributing via eigrp 1
  Last update from 120.100.123.3 on Serial0/0/0, 00:03:06 ago
  Routing Descriptor Blocks:
  * 120.100.123.3, from 120.100.123.3, 00:03:06 ago, via Serial0/0/0
    Route metric is 7289856, traffic share count is 1
    Total delay is 220000 microseconds, minimum bandwidth is 1544 Kbit
    Reliability 255/255, minimum MTU 1500 bytes
    Loading 1/255, Hops 1
```

```
R3(config-if)# interface Serial0/0/0
R3(config-if)# no ip next-hop-self eigrp 1
```

R1# **show ip route 150.101.2.0**

```
Routing entry for 150.101.2.0/24
  Known via "eigrp 1", distance 170, metric 7289856, type external
  Redistributing via eigrp 1
  Last update from 120.100.123.2 on Serial0/0/0, 00:00:24 ago
  Routing Descriptor Blocks:
  * 120.100.123.2, from 120.100.123.3, 00:00:24 ago, via Serial0/0/0
    Route metric is 7289856, traffic share count is 1
    Total delay is 220000 microseconds, minimum bandwidth is 1544 Kbit

    Reliability 255/255, minimum MTU 1500 bytes
    Loading 1/255, Hops 1
```

- Perform mutual redistribution of EIGRP and OSPF on R4 and R5. Use a metric of 5000 for redistributed routes into OSPF, which should appear as external Type 2 routes and the following K values for OSPF routes redistributed into EIGRP: 1544 20000 255 1 1500. (2 points)

This is an unambiguous redistribution question that sets the scene for the question that follows. Example 2-19 shows the required configuration on R4 and R5 with verification of external EIGRP received routes on R3. Because the metrics are identical on R4 and R5, there are multiple routes with load sharing potential. If you have configured this correctly, you have scored 2 points.

EXAMPLE 2-19 R4 and R5 Redistribution Configuration and Verification on R3

```
R4(config-router)# router ospf 1
R4(config-router)# redistribute eigrp 1 subnets
R4(config-router)# default-metric 5000
R4(config-router)# router eigrp 1
R4(config-router)# redistribute ospf 1
R4(config-router)# default-metric 1544 20000 255 1 1500

R5(config-router)# router ospf 1
R5(config-router)# redistribute eigrp 1 subnets
R5(config-router)# default-metric 5000
R5(config-router)# router eigrp 1
R5(config-router)# redistribute ospf 1
R5(config-router)# default-metric 1544 20000 255 1 1500

R3# show ip route eigrp
 150.101.0.0/24 is subnetted, 2 subnets
D       150.101.1.0 [90/2297856] via 120.100.123.1, 00:05:05, Serial0/0/0
 120.0.0.0/8 is variably subnetted, 20 subnets, 3 masks
D EX    120.100.9.1/32
         [170/6780416] via 120.100.34.5, 00:00:22, GigabitEthernet0/0
         [170/6780416] via 120.100.34.4, 00:00:22, GigabitEthernet0/0
D EX    120.100.8.1/32
         [170/6780416] via 120.100.34.5, 00:00:22, GigabitEthernet0/0
         [170/6780416] via 120.100.34.4, 00:00:22, GigabitEthernet0/0
D EX    120.100.10.1/32
         [170/6780416] via 120.100.34.5, 00:00:22, GigabitEthernet0/0
         [170/6780416] via 120.100.34.4, 00:00:22, GigabitEthernet0/0
```

```

D EX 120.100.5.1/32
      [170/6780416] via 120.100.34.4, 00:01:51, GigabitEthernet0/0
D    120.100.4.0/24
      [90/156160] via 120.100.34.4, 00:07:17, GigabitEthernet0/0
D    120.100.5.0/24
      [90/156160] via 120.100.34.5, 00:07:17, GigabitEthernet0/0
D EX 120.100.4.1/32
      [170/6780416] via 120.100.34.5, 00:00:23, GigabitEthernet0/0
D EX 120.100.7.1/32
      [170/6780416] via 120.100.34.5, 00:00:23, GigabitEthernet0/0
      [170/6780416] via 120.100.34.4, 00:00:23, GigabitEthernet0/0
D EX 120.100.6.1/32
      [170/6780416] via 120.100.34.5, 00:00:24, GigabitEthernet0/0
      [170/6780416] via 120.100.34.4, 00:00:24, GigabitEthernet0/0
D    120.100.0.0/16 [90/2172416] via 120.100.123.1, 00:05:07, Serial0/0/0
D    120.100.1.0/24 [90/2297856] via 120.100.123.1, 00:05:07, Serial0/0/0
D    120.100.2.0/24 [90/2297856] via 120.100.123.2, 00:05:07, Serial0/0/0
D EX 120.100.63.0/24
      [170/6780416] via 120.100.34.5, 00:00:24, GigabitEthernet0/0
      [170/6780416] via 120.100.34.4, 00:00:24, GigabitEthernet0/0
D EX 120.100.53.0/24
      [170/6780416] via 120.100.34.5, 00:00:24, GigabitEthernet0/0
      [170/6780416] via 120.100.34.4, 00:00:24, GigabitEthernet0/0
D EX 120.100.46.0/24
      [170/6780416] via 120.100.34.5, 00:00:24, GigabitEthernet0/0
      [170/6780416] via 120.100.34.4, 00:00:24, GigabitEthernet0/0
D    120.100.100.0/24 [90/2172416] via 120.100.123.1, 00:05:07, Serial0/0/0
D    120.100.200.0/24 [90/2172416] via 120.100.123.2, 00:05:08, Serial0/0/0

```

- R3 will have equal cost external EIGRP routes to the redistributed OSPF subnet 120.100.63.0/24 (VLAN 63). Configure only R3 to ensure that R3 routes via a next hop of R5 (120.100.34.5) for this destination subnet. If this route fails, the route advertised from R4 (120.100.34.4) should be used dynamically. (3 points)

Example 2-20 shows both routes for 120.100.63.0/24 received on R3 from R4 and R5; because all routers share a common media, the interface connecting to R4 or R5 cannot be modified on R3 because this would affect both routes. Similarly, an offset-list to manipulate delay would be of no use because you are permitted to configure only R3. You are therefore required to penalize the route received from R4 only to ensure the R5-generated route is preferred on R3. By configuring a route-map on R3 to match only the route-source of R4, you can increase the metric for the required route (120.100.63.0/24). This simply enables the original route received from R5 to take precedence. Example 2-20 shows the

required configuration and verification that the route is preferred via the R5, the topology table shows that the R4 route is also present and that R4 is effectively the feasible successor for this network on this router. If the route from R5 is withdrawn, the route from R5 would enter the routing table automatically. You will need a second permit statement on the route-map (permit 20) to enable all other routes inbound to R3 to enter unaltered. Example 2-20 also details the routing tables of each device to confirm redistribution from EIGRP into OSPF or vice versa. If you have configured this correctly, as shown in Example 2-20, you have scored 3 points.

EXAMPLE 2-20 R3 RIPv2 Redistribution Configuration and Verification

```
R3# show ip route 120.100.63.0
Routing entry for 120.100.63.0/24
  Known via "eigrp 1", distance 170, metric 6780416, type external
  Redistributing via eigrp 1
  Last update from 120.100.34.5 on GigabitEthernet0/0, 00:01:59 ago
  Routing Descriptor Blocks:
    120.100.34.5, from 120.100.34.5, 00:01:59 ago, via GigabitEthernet0/0
      Route metric is 6780416, traffic share count is 1
      Total delay is 200100 microseconds, minimum bandwidth is 1544 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 1
    * 120.100.34.4, from 120.100.34.4, 00:01:59 ago, via GigabitEthernet0/0
      Route metric is 6780416, traffic share count is 1
      Total delay is 200100 microseconds, minimum bandwidth is 1544 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 1

R3(config)# access-list 1 permit 120.100.34.4
R3(config)# access-list 2 permit 120.100.63.0
R3(config)# router eigrp 1
R3(config-router)# distribute-list route-map PENALISE-VLAN63 in GigabitEthernet0/0
R3(config-router)# exit
R3(config)# route-map PENALISE-VLAN63 permit 10
R3(config-route-map)# match ip address 2
R3(config-route-map)# match ip route-source 1
R3(config-route-map)# set metric +500000
R3(config-route-map)# route-map PENALISE-VLAN63 permit 20

R3# show ip route 120.100.63.0
```

```
Routing entry for 120.100.63.0/24
  Known via "eigrp 1", distance 170, metric 6780416, type external
  Redistributing via eigrp 1
  Last update from 120.100.34.5 on GigabitEthernet0/0, 00:00:21 ago
  Routing Descriptor Blocks:
  * 120.100.34.5, from 120.100.34.5, 00:00:21 ago, via GigabitEthernet0/0
    Route metric is 6780416, traffic share count is 1
    Total delay is 200100 microseconds, minimum bandwidth is 1544 Kbit
    Reliability 255/255, minimum MTU 1500 bytes
    Loading 1/255, Hops 1
```

R3# **show ip eigrp topology 120.100.63.0 255.255.255.0**

```
IP-EIGRP (AS 1): Topology entry for 120.100.63.0/24
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 6780416
  Routing Descriptor Blocks:
  120.100.34.5 (GigabitEthernet0/0), from 120.100.34.5, Send flag is 0x0
    Composite metric is (6780416/6777856), Route is External
    Vector metric:
      Minimum bandwidth is 1544 Kbit
  Total delay is 200100 microseconds
    Reliability is 255/255
    Load is 1/255
    Minimum MTU is 1500
    Hop count is 1
  External data:
    Originating router is 120.100.5.1
    AS number of route is 1
    External protocol is OSPF, external metric is 2
    Administrator tag is 0 (0x00000000)
  120.100.34.4 (GigabitEthernet0/0), from 120.100.34.4, Send flag is 0x0
    Composite metric is (128000000/6777856), Route is External
    Vector metric:
      Minimum bandwidth is 20 Kbit
      Total delay is 0 microseconds
      Reliability is 0/255
      Load is 0/255
      Minimum MTU is 0
      Hop count is 1
  External data:
    Originating router is 120.100.4.1
```

NOTE

The full IP routing tables of each device are provided within the accompanying configurations to verify your redistributed routes.

```
AS number of route is 1
External protocol is OSPF, external metric is 2
Administrator tag is 0 (0x00000000)
```

Section 3: BGP (15 Points)

- Configure BGP peering per Figure 2-9 as follows: iBGP R1-R3, R2-R3, R4-R6, R4-SW2, R5-Sw1, R5-sw3, eBGP R3-R4, R3-R5, Sw4-Sw3, R6-Sw4. Use Loopback interfaces to peer on all routers with the exception of peering between R3-R4 and R3-R5. Do not use the command **ebgp-multihop** within your configurations. (3 points)

The restrictions within the iBGP peering require you to configure R3, R4, and R5 as route reflectors within their own AS. Auto summarization is disabled to ensure BGP does not summarize routes, and synchronization is disabled because the IGP will not be synchronized to BGP within this lab. The question doesn't dictate that you must configure peer groups, but it is considered good practice when you have more than one peer with a similar peering configuration. The question does, however, dictate that you must not use **ebgp-multihop**. This feature would of course be required for the peering from AS400 to AS300 and AS400 to AS200 because Loopback interfaces are used for the external peering, here unlike AS100 to AS200 and AS300, which peer from connected interfaces. Without **ebgp-multihop** the peering fails in and outbound from AS400. The only way to fix this is to use a feature that disables connection verification to establish an eBGP peering session with a single-hop peer that uses a Loopback interface. Use of the command **neighbor disable-connected-check** on R6, Sw3, and Sw4 for the required peering allows the peering to be formed successfully. Example 2-21 shows the basic peering configuration for BGP, the eBGP failure condition observed on peering to and from AS400, and the required configuration to rectify the condition. If you have configured this correctly, you have scored 3 points.

EXAMPLE 2-21 BGP Peering Configuration and Verification

```
R1(config)# router bgp 100
R1(config-router)# no auto-summary
R1(config-router)# no synchronization
R1(config-router)# neighbor 120.100.3.1 remote-as 100
R1(config-router)# neighbor 120.100.3.1 update-source Loopback0

R2(config)# router bgp 100
R2(config-router)# no auto-summary
R2(config-router)# no synchronization
R2(config-router)# neighbor 120.100.3.1 remote-as 100
R2(config-router)# neighbor 120.100.3.1 update-source Loopback0
```

```
R3(config)# router bgp 100
R3(config-router)# no auto-summary
R3(config-router)# no synchronization
R3(config-router)# neighbor AS100 peer-group
R3(config-router)# neighbor AS100 remote-as 100
R3(config-router)# neighbor AS100 update-source Loopback0
R3(config-router)# neighbor 120.100.1.1 peer-group AS100
R3(config-router)# neighbor 120.100.2.1 peer-group AS100
R3(config-router)# neighbor AS100 route-reflector-client
R3(config-router)# neighbor 120.100.34.4 remote-as 200
R3(config-router)# neighbor 120.100.34.5 remote-as 300
```

```
R4(config)# router bgp 200
R4(config-router)# router bgp 200
R4(config-router)# no auto-summary
R4(config-router)# no synchronization
R4(config-router)# neighbor AS200 peer-group
R4(config-router)# neighbor AS200 remote-as 200
R4(config-router)# neighbor AS200 update-source Loopback0
R4(config-router)# neighbor AS200 route-reflector-client
R4(config-router)# neighbor 120.100.6.1 peer-group AS200
R4(config-router)# neighbor 120.100.8.1 peer-group AS200
R4(config-router)# neighbor 120.100.34.3 remote-as 100
```

```
R5(config)# router bgp 300
R5(config-router)# no auto-summary
R5(config-router)# no synchronization
R5(config-router)# neighbor AS300 peer-group
R5(config-router)# neighbor AS300 remote-as 300
R5(config-router)# neighbor AS300 update-source Loopback0
R5(config-router)# neighbor AS300 route-reflector-client
R5(config-router)# neighbor 120.100.7.1 peer-group AS300
R5(config-router)# neighbor 120.100.9.1 peer-group AS300
R5(config-router)# neighbor 120.100.34.3 remote-as 100
```

```
R6(config)# router bgp 200
R6(config-router)# no auto-summary
```

```
R6(config-router)# no synchronization
R6(config-router)# neighbor 120.100.4.1 remote-as 200
R6(config-router)# neighbor 120.100.4.1 update-source Loopback0
R6(config-router)# neighbor 120.100.10.1 remote-as 400
R6(config-router)# neighbor 120.100.10.1 update-source Loopback0
```

```
SW1(config)# router bgp 300
SW1(config-router)# no auto-summary
SW1(config-router)# no synchronization
SW1(config-router)# neighbor 120.100.5.1 remote-as 300
SW1(config-router)# neighbor 120.100.5.1 update-source Loopback0
```

```
SW2(config)# router bgp 200
SW2(config-router)# no auto-summary
SW2(config-router)# no synchronization
SW2(config-router)# neighbor 120.100.4.1 remote-as 200
SW2(config-router)# neighbor 120.100.4.1 update-source Loopback0
```

```
SW3(config)# router bgp 300
SW3(config-router)# no auto-summary
SW3(config-router)# no synchronization
SW3(config-router)# neighbor 120.100.5.1 remote-as 300
SW3(config-router)# neighbor 120.100.5.1 update-source Loopback0
SW3(config-router)# neighbor 120.100.10.1 remote-as 400
SW3(config-router)# neighbor 120.100.10.1 update-source Loopback0
```

```
SW4(config)# router bgp 400
SW4(config-router)# no auto-summary
SW4(config-router)# no synchronization
SW4(config-router)# neighbor 120.100.6.1 remote-as 200
SW4(config-router)# neighbor 120.100.6.1 update-source Loopback0
SW4(config-router)# neighbor 120.100.9.1 remote-as 300
SW4(config-router)# neighbor 120.100.9.1 update-source Loopback0
```

```
SW4# sh ip bgp neigh 120.100.6.1 | include External
External BGP neighbor not directly connected.
SW4# show ip bgp neighbors 120.100.9.1 | include External
External BGP neighbor not directly connected.
```

```
SW4#

SW4# sh ip bgp neighbors 120.100.6.1 | include active
No active TCP connection
SW4# sh ip bgp neighbors 120.100.9.1 | include active
No active TCP connection

SW4(config-router)# neighbor 120.100.6.1 disable-connected-check
SW4(config-router)# neighbor 120.100.9.1 disable-connected-check

R6(config-router)# neighbor 120.100.10.1 disable-connected-check

SW3(config-router)# neighbor 120.100.10.1 disable-connected-check

SW4# show ip bgp neighbors 120.100.6.1 | include Established
BGP state = Established, up for 00:02:01
SW4# show ip bgp neighbors 120.100.9.1 | include Established
BGP state = Established, up for 00:02:05
```

You will also find peering issues between R1 and R3. Example 2-22 shows the routers are informing each other they have an incorrect BGP identifier. This is simply because both routers have identical Loopback interface address of 200.200.200.200, which is used as the BGP identifier. By changing the ID of one router the peering is established. It doesn't matter what you change the ID to, but it needs to be unique; as such, the Loopback 0 interface would be a good choice. No extra points for this task because this is part of the original peering.

EXAMPLE 2-22 R1 and R3 Peering Issue Configuration and Verification

```
R1# * 19:30:13.287: %BGP-3-NOTIFICATION: sent to neighbor 120.100.3.1 2/3 (BGP
identifier wrong) 4 bytes C8C8C8C8

R3# * 19:25:30.043: %BGP-3-NOTIFICATION: received from neighbor 120.100.1.1 2/
3 (BGP identifier wrong) 4 bytes C8C8C8C8

R1# show ip bgp summary | include identifier
BGP router identifier 200.200.200.200, local AS number 100

R3# show ip bgp summary | include identifier
BGP router identifier 200.200.200.200, local AS number 100
```

```
R1(config-router)#  bgp router-id 120.100.1.1  
*19:34:45.467:  %BGP-5-ADJCHANGE:  neighbor 120.100.3.1 Up
```

- Routers R1 and R2 in AS100 should be made to passively accept only BGP sessions. R3 should be configured to actively create only BGP sessions to R1 and R2 within AS100. (3 points)

A BGP speaker by default will attempt to open a session on TCP port 179 with a configured peer, because such a normal peering arrangement will see two sessions being established to build a successful neighbor relationship. This behavior can be modified to effectively allow sessions to be established only either inbound or outbound. The solution to the question is achieved by configuring the **neighbor transport connection-mode** to passive (only inbound connections will be established) on R1 and R2 and active (only outbound sessions will be established) on R3. You must manually activate each neighbor on each router for the solution to work effectively. If you have configured this correctly, as shown in Example 2-23, you have scored 3 points. Consider using the **show ip bgp summary** command to verify your configura-

XAMPLE 2-23 R1, R2 and R3 Connection-mode Configuration

```
R1(config)# router bgp 100  
R1(config-router)# neighbor 120.100.3.1 transport connection-mode passive
```

```
R1(config-router)# neighbor 120.100.3.1 activate
```

```
R2(config)# router bgp 100  
R2(config-router)# neighbor 120.100.3.1 transport connection-mode passive  
R3(config)# router bgp 100  
R2(config-router)# neighbor 120.100.3.1 activate  
R3(config-router)# neighbor AS100 transport connection-mode active  
R3(config-router)# neighbor 120.100.1.1 activate  
R3(config-router)# neighbor 120.100.2.1 activate
```

- Configure the following Loopback interfaces on R3 and Sw4; advertise these networks into BGP using the network command: (2 points)

R3 – Loopback interface 5 (152.100.100.1/24) Sw4 – Loopback interface 5 (152.200.32.1/24)
Sw4 – Loopback interface 6 (152.200.33.1/24)

Sw4 – Loopback interface 7 (152.200.34.1/24)

Sw4 – Loopback interface 8 (152.200.35.1/24)

A simple question that creates BGP routes for the following task. If you have configured this correctly, as shown in Example 2-24, you have scored 2 points.

EXAMPLE 2-24 R3 and Sw4 Network Advertisement Configuration and Verification

```
R3(config)# interface Loopback5
R3(config-if)# ip address 152.100.100.1 255.255.255.0
R3(config-if)# router bgp 100
R3(config-router)# network 152.100.100.0 mask 255.255.255.0

SW4(config)# interface Loopback5
SW4(config-if)# ip address 152.200.32.1 255.255.255.0
SW4(config-if)# interface Loopback6
SW4(config-if)# ip address 152.200.33.1 255.255.255.0
SW4(config-if)# interface Loopback7
SW4(config-if)# ip address 152.200.34.1 255.255.255.0
SW4(config-if)# interface Loopback8
SW4(config-if)# ip address 152.200.35.1 255.255.255.0
SW4(config-if)# router bgp 400
SW4(config-router)# network 152.200.32.0 mask 255.255.255.0
SW4(config-router)# network 152.200.33.0 mask 255.255.255.0
SW4(config-router)# network 152.200.34.0 mask 255.255.255.0
SW4(config-router)# network 152.200.35.0 mask 255.255.255.0

R3# show ip bgp
BGP table version is 10, local router ID is 200.200.200.200
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop           Metric LocPrf  Weight  Path
*> 152.100.100.0/24  0.0.0.0             0         32768   i
*   152.200.32.0/24  120.100.34.4        0          200 400  i
*>                   120.100.34.5        0          300 400  i
*   152.200.33.0/24  120.100.34.4        0          200 400  i
```

```

*>                120.100.34.5                0 300 400 i
* 152.200.34.0/24 120.100.34.4                0 200 400 i
*>                120.100.34.5                0 300 400 i
* 152.200.35.0/24 120.100.34.4                0 200 400 i
*>                120.100.34.5                0 300 400 i

```

- Configure R3 to inform R4 that it does not want to receive routes advertised from Sw4 for networks 152.200.33.0/24, 152.200.34.0/24 and 152.200.35.0/24. Achieve this in such a manner that R4 does not actually advertise these routes toward R3. You may also configure R4. (4 points)

BGP has a Prefix-Based Outbound Route Filtering (ORF) mechanism that can send and receive capabilities to minimize BGP updates sent between BGP peers. Advertisement of ORF capability indicates that a peer will accept a prefix-list from a neighbor and apply the prefix-list received from a neighbor locally to avoid the unnecessary sending of routes that would be blocked by the receiver anyway. R3 is therefore configured with a prefix-list that blocks the required routes generated from Sw4, which is sent via ORF to R4. R4 is configured to receive this prefix-list via ORF, and the routes are blocked outbound at R4. Example 2-25 shows the required ORF and prefix-list filtering with the resulting outbound advertisement on R4. The BGP table on R3 is also displayed showing the routes are no longer being received from R4 and solely from R5. If you have configured this correctly, as shown in Example 2-25, you have scored 4 points.

EXAMPLE 2-25 BGP ORF Configuration and Verification

```

R3(config)#  router bgp 100
R3(config-router)#  neighbor 120.100.34.4 capability orf prefix-list send
R3(config-router)#  neighbor 120.100.34.4 prefix-list FILTER in
R3(config)#  ip prefix-list FILTER seq 5 deny 152.200.33.0/24
R3(config)#  ip prefix-list FILTER seq 10 deny 152.200.34.0/24
R3(config)#  ip prefix-list FILTER seq 15 deny 152.200.35.0/24
R3(config)#  ip prefix-list FILTER seq 20 permit 0.0.0.0 le 32

R4(config)#  router bgp 200
R4(config-router)#  neighbor 120.100.34.3 capability orf prefix-list receive
R4(config-router)#  exit
R4(config)#  exit

R4#  show ip bgp neighbors 120.100.34.3 advertised-routes
BGP table version is 17, local router ID is 120.100.4.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

```

```

      Network          Next Hop          Metric LocPrf Weight Path
*> 152.200.32.0/24    120.100.10.1          0     100     0 400 i

```

```
Total number of prefixes 1
```

```
R3# clear ip bgp *
```

```
R3# show ip bgp
```

```
BGP table version is 6, local router ID is 200.200.200.200
```

```
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale
```

```
Origin codes: i - IGP, e - EGP, ? - incomplete
```

```

      Network          Next Hop          Metric LocPrf Weight Path
*> 152.100.100.0/24    0.0.0.0              0           32768 i
*> 152.200.32.0/24    120.100.34.4          0     200 400 i
*                   120.100.34.5          0     300 400 i
*> 152.200.33.0/24    120.100.34.5          0     300 400 i
*> 152.200.34.0/24    120.100.34.5          0     300 400 i
*> 152.200.35.0/24    120.100.34.5          0     300 400 i

```

- Configure a route-map on R5 that prepends its local AS 2 an additional 2 times for network 152.200.32.0/24 when advertised to R3. The route-map may contain multiple permit statements but only one prepend is permitted per line. (3 points)

A simple AS path prepend question, or so it seems. Normally you would prepend the same AS number multiple times within the same permit statement, but the question restricts this so you are forced to use multiple permit statements with the same AS prepend statement. Example 2-26 shows the route 152.200.32.0/24 as received initially on R3 from R5 with an AS path of 300-400. After configuration of the route-map to prepend the route on R5 twice, the network is received on R3 with an AS path of 300-300-400. This might look like the route has indeed been prepended twice, but the question requests an “additional” two times; in fact, the route has been prepended only once. The problem is that the route-map **permit 10** statement on R3 has been executed, and the route-map will then not evaluate any additional route map entries and simply drops out, so the permit 20 statement is never actually executed. By configuring a **continue 20** statement within the **permit 10** line, the router is forced to evaluate the permit 20 line. Rather than dropping out of the route-map after successful execution of the **permit 10** statement, the final verification within Example 2-26 shows the route received on R3 with successful prepend applied by R5. If you have configured this correctly, as shown in Example 2-26, you have scored 3 points.

EXAMPLE 2-26 R5 Prepend Configuration and Verification

```
R3# show ip bgp
BGP table version is 6, local router ID is 200.200.200.200
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
```

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 152.100.100.0/24	0.0.0.0	0		32768	i
*> 152.200.32.0/24	120.100.34.4			0 200	400 i
*	120.100.34.5			0 300	400 i
*> 152.200.33.0/24	120.100.34.5			0 300	400 i
*> 152.200.34.0/24	120.100.34.5			0 300	400 i
*> 152.200.35.0/24	120.100.34.5			0 300	400 i

```
R5(config)# router bgp 300
R5(config-router)# neighbor 120.100.34.3 route-map PREPEND out
R5(config-router)# exit
R5(config)# access-list 1 permit 152.200.32.0
R5(config)# route-map PREPEND permit 10
R5(config-route-map)# match ip address 1
R5(config-route-map)# set as-path prepend 300
R5(config-route-map)# route-map PREPEND permit 20
R5(config-route-map)# match ip address 1
R5(config-route-map)# set as-path prepend 300
R5(config-route-map)# route-map PREPEND permit 30
```

```
R3# show ip bgp
BGP table version is 6, local router ID is 200.200.200.200
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
```

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 152.100.100.0/24	0.0.0.0	0		32768	i
*> 152.200.32.0/24	120.100.34.4			0 200	400 i
*	120.100.34.5			0 300 300	400 i
*> 152.200.33.0/24	120.100.34.5			0 300	400 i
*> 152.200.34.0/24	120.100.34.5			0 300	400 i

```
*> 152.200.35.0/24 120.100.34.5 0 300 400 i

R5(config)# route-map PREPEND permit 10
R5(config-route-map)# continue 20

R3# clear ip bgp *
R3# show ip bgp
BGP table version is 6, local router ID is 200.200.200.200
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
*> 152.100.100.0/24 0.0.0.0            0         32768  i
*> 152.200.32.0/24 120.100.34.4      0         200 400  i
*                   120.100.34.5      0         300 300 300 400  i
*> 152.200.33.0/24 120.100.34.5      0         300 400  i
*> 152.200.34.0/24 120.100.34.5      0         300 400  i
*> 152.200.35.0/24 120.100.34.5      0         300 400  i
```

Section 4: IPv6 (12 Points)

- Configure IPv6 addresses on your network as follows:

```
2007:C15:C0:10::1/64 - R1 Gi0/0
2007:C15:C0:11::1/64 - R1 tunnel0
2007:C15:C0:11::3/64 - R3 tunnel0
2007:C15:C0:12::2/64 - R2 tunnel0
2007:C15:C0:12::3/64 - R3 tunnel1
2007:C15:C0:13::2/64 - R2 fe0/1
2007:C15:C0:14::3/64 - R3 Gi0/0
2007:C15:C0:14::4/64 - R4 Gi0/0
```

2007:C15:C0:14::5/64 – R5 Gi0/0

2007:C15:C0:15::5/64 – R4 Gi0/1

2007:C15:C0:15::6/64 – R6 Gi0/0

The prerequisite to the following questions is configuration of the IPv6 addresses and tunnel interfaces. You should test your IPv6 connectivity post configuration to ensure you are ready to progress to the routing questions. You will not require Frame-Relay maps to achieve connectivity because tunneling is required rather than IPv6 directly configured under the serial interfaces on R1, R2, and R3. Example 2-27 shows the initial IPv6 configuration; tunnel specifics are provided in later questions, so just creating the tunnel interfaces and configuring an IPv6 address is required at this point. No points are on offer here for this task, unfortunately. Consider using the **show ipv6 interfaces brief** command for a quick check of your interface configuration.

EXAMPLE 2-27 IPv6 Initial Configuration

```
R1(config)# ipv6 unicast-routing
R1(config)# interface GigabitEthernet0/1
R1(config-if)# ipv6 address 2007:C15:C0:10::1/64
R1(config-if)# interface tunnel0
R1(config-if)# ipv6 address 2007:C15:C0:11::1/64

R2(config)# ipv6 unicast-routing
R2(config)# interface FastEthernet 0/1
R2(config-if)# ipv6 address 2007:C15:C0:13::2/64
R2(config-if)# interface tunnel0
R2(config-if)# ipv6 address 2007:C15:C0:12::2/64

R3(config)# ipv6 unicast-routing
R3(config)# int GigabitEthernet0/0
R3(config-if)# ipv6 address 2007:C15:C0:14::3/64
R3(config-if)# interface tunnel0
R3(config-if)# ipv6 address 2007:C15:C0:11::3/64
R3(config-if)# interface tunnel1
R3(config-if)# ipv6 address 2007:C15:C0:12::3/64

R4(config)# ipv6 unicast-routing
R4(config)# interface GigabitEthernet0/0
```

```
R4(config-if)# ipv6 address 2007:C15:C0:14::4/64
R4(config-if)# interface GigabitEthernet0/1
R4(config-if)# ipv6 address 2007:C15:C0:15::4/64

R5(config)# ipv6 unicast-routing
R5(config)# interface GigabitEthernet0/0
R5(config-if)# ipv6 address 2007:C15:C0:14::5/64

R6(config)# ipv6 unicast-routing
R6(config)# interface GigabitEthernet0/0
R6(config-if)# ipv6 address 2007:C15:C0:15::6/64
```

Section 4.1: EIGRPv6

- Configure EIGRPv6 between R1, R2, and R3. EIGRPv6 should be enabled on the Ethernet interfaces of R1 and R2 and on all tunnel interfaces of R1, R2, and R3. Build your tunnels using **ipv6ip** mode—use an AS number of 6 on all required interfaces. (2 points)

This is a straightforward EIGRPv6 configuration that requires the AS number of 6 applied to the required interfaces. The tunnel mode information is supplied within this question of **ipv6ip** for a manually configured IPv6 tunnel. One thing to remember with EIGRPv6 is that you need to start the process with a **no shut** command within the routing process. If you have configured this correctly, as shown in Example 2-28, you have scored 2 points.

EXAMPLE 2-28 EIGRPv6 Configuration and Verification

```
R1(config)# interface GigabitEthernet 0/1
R1(config-if)# ipv6 eigrp 6
R1(config-if)# interface Tunnel0
R1(config-if)# ipv6 eigrp 6
R1(config-if)# tunnel source Serial0/0/0
R1(config-if)# tunnel destination 120.100.123.3
R1(config-if)# tunnel mode ipv6ip
R1(config-if)# ipv6 router eigrp 6
R1(config-router)# no shutdown

R2(config)# interface FastEthernet 0/1
R2(config-if)# ipv6 eigrp 6
R2(config-if)# interface Tunnel0
```

```
R2(config-if)# ipv6 eigrp 6
R2(config-if)# tunnel source Serial0/0
R2(config-if)# tunnel destination 120.100.123.3
R2(config-if)# tunnel mode ipv6ip
R2(config-if)# ipv6 router eigrp 6
R2(config-router)# no shutdown
```

```
R3(config)# interface Tunnel0
R3(config-if)# ipv6 eigrp 6
R3(config-if)# tunnel source Serial0/0/0
R3(config-if)# tunnel destination 120.100.123.1
R3(config-if)# tunnel mode ipv6ip
R3(config-if)# interface Tunnel1
R3(config-if)# ipv6 eigrp 6
R3(config-if)# tunnel source Serial0/0/0
R3(config-if)# tunnel destination 120.100.123.2
R3(config-if)# tunnel mode ipv6ip
R3(config-if)# ipv6 router eigrp 6
R3(config-router)# no shutdown
```

R1# show ipv6 route eigrp

```
IPv6 Routing Table - 8 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B -
       BGP U - Per-user Static route, M - MIPv6
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
       D - EIGRP, EX - EIGRP external
D 2007:C15:C0:12::/64 [90/310044416]
   via FE80::7864:7B03, Tunnel0
D 2007:C15:C0:13::/64 [90/310070016]
   via FE80::7864:7B03, Tunnel0
```

R2# show ipv6 route eigrp

```
IPv6 Routing Table - 8 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B -
       BGP U - Per-user Static route, M - MIPv6
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
```

```

        ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
        D - EIGRP, EX - EIGRP external
D 2007:C15:C0:10::/64 [90/310070016]
    via FE80::7864:7C03, Tunnel0
D 2007:C15:C0:11::/64 [90/310044416]
    via FE80::7864:7C03, Tunnel0

R3# show ipv6 route eigrp
IPv6 Routing Table - 9 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B -
        BGP U - Per-user Static route, M - MIPv6
        I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
        O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
        ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
        D - EIGRP, EX - EIGRP external
D 2007:C15:C0:10::/64 [90/297270016]
    via FE80::7864:7B01, Tunnel0
D 2007:C15:C0:13::/64 [90/297270016]
    via FE80::7864:7C02, Tunnel1

```

Section 4.2: OSPFv3

- Configure OSPFv3 per Figure 2-11; use an OSPFv3 process of 1 on each router. (2 points)

Use vanilla OSPFv3 configuration between R3, R4, R5, and R6. If you have configured this correctly, as shown in Example 2-29, you have scored 2 points.

EXAMPLE 2-29 OSPFv3 Configuration and Verification

```

R3(config)# interface GigabitEthernet 0/0
R3(config-if)# ipv6 ospf 1 area 0

R4(config)# interface GigabitEthernet0/0
R4(config-if)# ipv6 ospf 1 area 0
R4(config-if)# interface GigabitEthernet0/1
R4(config-if)# ipv6 ospf 1 area 1

R5(config)# interface GigabitEthernet0/0
R5(config-if)# ipv6 ospf 1 area 0

```

```
R6(config)# interface GigabitEthernet0/0
```

```
R6(config-if)# ipv6 ospf 1 area 1
```

```
R3# show ipv6 route ospf
```

```
IPv6 Routing Table - 11 entries
```

```
Codes: C - Connected, L - Local, S - Static, R - RIP, B -
```

```
       BGP U - Per-user Static route
```

```
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
```

```
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
```

```
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
```

```
       D - EIGRP, EX - EIGRP external
```

```
OI 2007:C15:C0:15::/64 [110/2]
```

```
   via FE80::213:C3FF:FE7B:E4A0, GigabitEthernet0/0
```

```
R5# show ipv6 route ospf
```

```
IPv6 Routing Table - 5 entries
```

```
       U - Per-user Static route
```

```
Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP
```

```
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
```

```
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
```

```
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
```

```
       D - EIGRP, EX - EIGRP external
```

```
OI 2007:C15:C0:15::/64 [110/2]
```

```
   via FE80::213:C3FF:FE7B:E4A0, GigabitEthernet0/0
```

```
R6# show ipv6 route ospf
```

```
IPv6 Routing Table - 5 entries
```

```
Codes: C - Connected, L - Local, S - Static, R - RIP, B -
```

```
       BGP U - Per-user Static route
```

```
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
```

```
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
```

```
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
```

```
       D - EIGRP, EX - EIGRP external
```

```
OI 2007:C15:C0:14::/64 [110/2]
```

```
   via FE80::213:C3FF:FE7B:E4A1, GigabitEthernet0/0
```

NOTE

The OSPFv3 routing table of R4 is not shown in Example 2-29 because this router physically connects to each IPv6 network and as such will not discover any OSPFv3 dynamic routes at this point in time.

- Configure Area 1 with IPsec authentication, use Message Digest 5, a Security Policy Index of 500, and a key of **DECODECC1E0DDBA11B0BB0BBEDB00B00**. (2 points)

Authentication is required on R4 and R6 because they both belong to Area 1. The question explicitly states the specific parameters required, and you shouldn't encounter any issues unless you incorrectly enter one of the keys. At 32 Hex characters long, this could easily be done while under a time constraint. If you have configured this correctly, as shown in Example 2-30, you have scored 2 points.

EXAMPLE 2-30 Area 1 Authentication Configuration

```
R4(config)# ipv6 router ospf 1
R4(config-router)# area 1 authentication ipsec spi 500 md5 DEC0DECC1E0DDBA11B0BB0BBEDB00B00

R6(config)# ipv6 router ospf 1
R6(config-router)# area 1 authentication ipsec spi 500 md5 DEC0DECC1E0DDBA11B0BB0BBEDB00B00
```

- Ensure the area router in Area 1 receives the following route; you may configure R4 to achieve this: (2 points) OI
2007::/16 [110/2]
via XXXX::XXXX:XXXX:XXXX:XXXX, GigabitEthernet0/0

The only area router within Area 1 is R6. R4 is the area border router within this area. OI within the routing table is an OSPF Interarea route, so this route must be generated from another area. Because Area 0 is the only other area within the OSPFv3 network, the route must be generated from this area as opposed to a redistributed route, which would show as an external route. A summary route generated on the area border Router R4 of 2007::/16 within area 0 will provide the required route to be received on R6. If you have configured this correctly, as shown in Example 2-31, you have scored 2 points.

EXAMPLE 2-31 OSPFv3 Configuration and Verification

```
R4(config)# ipv6 router ospf 1
R4(config-rtr)# area 0 range 2007::/16

R6# show ipv6 route ospf | include OI
OI 2007::/16 [110/2]
   via FE80::213:C3FF:FE7B:E4A1, GigabitEthernet0/0
```

Section 4.3: Redistribution

- Redistribute EIGRPv6 into OSPFv3 on R3. Redistributed EIGRPv6 routes should have a metric of 5000 associated with them, regardless of which area they are seen in within the OSPFv3 network. (2 points)

A one-way redistribution of EIGRPv6 to OSPFv3 is required on R3. The default redistribution behavior ensures that external routes are advertised as external Type 2, which have a fixed cost associated with them regardless of which area or location of the OSPFv3 network they are seen in. You simply require the metric set to 5000 on the OSPFv3 process. You need to remember to advertise connected routes also; otherwise, the OSPFv3 network will not see the directly connected tunnel interfaces on R3. If you have configured this correctly, as shown in Example 2-32, you have scored 2 points.

EXAMPLE 2-32 R3 IPv6 Redistribution Configuration and Verification

```
R3(config)# ipv6 router ospf 1
R3(config-rtr)# redistribute eigrp 6 include-connected metric 5000

R4# show ipv6 route ospf
IPv6 Routing Table - 11 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B -
       BGP U - Per-user Static route
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
       D - EIGRP, EX - EIGRP external
O    2007::/16 [110/0]
    via ::, Null0
OE2  2007:C15:C0:10::/64 [110/5000]
    via FE80::214:6AFF:FEFC:7390, GigabitEthernet0/0
OE2  2007:C15:C0:11::/64 [110/5000]
    via FE80::214:6AFF:FEFC:7390, GigabitEthernet0/0
OE2  2007:C15:C0:12::/64 [110/5000]
    via FE80::214:6AFF:FEFC:7390, GigabitEthernet0/0
OE2  2007:C15:C0:13::/64 [110/5000]
    via FE80::214:6AFF:FEFC:7390, GigabitEthernet0/0
```

- Configure R3 so that both R1 and R2 have the following IPv6 EIGRPv6 route in place; do not redistribute OSPF into EIGRPv6 to achieve this, and ensure all routers have full visibility. (2 points)

```
D 2007::/16 [90/XXXXXXXXXX]
   via XXXX::XXXX:XXXX:XXXX:XXXX, Tunnel0
```

You should have noticed in the previous question that mutual redistribution was not required; as such, the EIGRPv6 network would not have reachability of the OSPFv3 network. This question ensures the RIPng network sends traffic to R3 for the summarized network of 2007::/16. Because you are not permitted to redistribute OSPFv3 with a summary address, you need to configure EIGRPv6 summarization on the tunnel interfaces on R3 toward R1 and R2; this will provide the correct route and hop count as per the question. Example 2-33 shows the required configuration and verification of the route, in addition to ICMP reachability to the remote OSPFv3 Area 1 network on R6. This test clearly demonstrates full end-to-end reachability from EIGRPv6 to OSPFv3. If you have configured this correctly, as shown in Example 2-33, you have scored 2 points.

EXAMPLE 2-33 R3 Ipv6 Summarization Configuration and Verification

```
R3(config)# interface tunnel 0
R3(config-if)# ipv6 summary-address eigrp 6 2007::/16
R3(config-if)# interface tunnel 1
R3(config-if)# ipv6 summary-address eigrp 6 2007::/16

R1# show ipv6 route eigrp
IPv6 Routing Table - 6 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B -
       BGP U - Per-user Static route, M - MIPv6
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
       D - EIGRP, EX - EIGRP external
D 2007::/16 [90/310044416]
  via FE80::7864:7B03, Tunnel0
R1# ping ipv6 2007:C15:C0:15::6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2007:C15:C0:15::6, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 4/7/8 ms

R2# show ipv6 route eigrp
```

```

IPv6 Routing Table - 6 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B -
       BGP U - Per-user Static route, M - MIPv6
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
       D - EIGRP, EX - EIGRP external
D 2007::/16 [90/310044416]
  via FE80::7864:7C03, Tunnel0

```

```
R2# ping ipv6 2007:C15:C0:15::6
```

```

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2007:C15:C0:15::6, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 4/7/8 ms

```

Section 5: QoS (6 Points)

- Two IP videoconferencing units are to be installed onto Switch 2 ports FastEthernet 0/15 and 0/16 on VLAN 200. The devices use TCP Ports 3230–3231 and UDP Ports 3230–3235, and this traffic is unmarked from the devices because it enters the switch. Configure Switch 2 to assign a DSCP value of AF41 to video traffic from both of these devices. Ensure that the switch ports assigned to the devices do not participate in the usual spanning-tree checks, cannot form trunk links, and cannot be configured as Etherchannels. (3 points)

This is a DSCP coloring of application traffic question. The TCP and UDP port information is provided so access-lists matching these ports within a class-map are required for identification of the video traffic, and a policy-map colors the traffic to a DSCP value of 41. The overall QoS service-policy is applied to the videoconferencing ports of FastEthernet 0/15 and 0/16 on Switch 2. The ports are required to be set to VLAN 200 with spanning-tree checks disabled, and trunking and channeling disabled using the command **switchport host**. The ports can also be explicitly configured to disable each feature individually but the **switchport host** command does all this for you. If you have configured this correctly, as shown in Example 2-34, you have scored 3 points. Use the **show policy-map** command to verify your configuration.

EXAMPLE 2-34 OSPFv3 Configuration

```

SW2(config)# interface range fastEthernet 0/15-16
SW2(config-if-range)# switchport access vlan 200
SW2(config-if-range)# switchport host

```

```
SW2(config-if-range)# exit
SW2(config)# mls qos
SW2(config)# class-map VIDEO
SW2(config-cmap)# match access-group 100
SW2(config-cmap)# exit
SW2(config)# access-list 100 permit tcp any any range 3230 3231
SW2(config)# access-list 100 permit udp any any range 3230 3235
SW2(config)# policy-map VIDEO-MARK
SW2(config-pmap)# class VIDEO
SW2(config-pmap-c)# set dscp AF41
SW2(config-pmap-c)# exit
SW2(config)# interface range fastEthernet 0/15-16
SW2(config-if-range)# service-policy input VIDEO-MARK
```

- Configure R2 to assign a strict priority queue with a 40-percent reservation of the WAN bandwidth for the video-conferencing traffic in the previous question. Maximize the available bandwidth by ensuring the RTP headers within the video stream are compressed. The remainder of the bandwidth should be guaranteed for a default queue with WRED enabled. Assume the full line rate of 1.544 Mbps as the available WAN bandwidth, and ensure the complete bandwidth is utilized by both queues. (3 points)

Following from the previous question, R2 is required to provide QoS on the Frame-Relay link. A class-map matches the precolored video traffic of DSCP 41; a policy-map is then required to call the class-map and assign a strict 40 percent priority queue with the command **priority percent 40**. RTP compression is configured within the policy-map for the video traffic. The default queue has a guaranteed bandwidth reservation with the command **bandwidth percent 60**, and WRED is enabled within this queue. Both queues can use the full bandwidth of the WAN link only if the command **max-reserved-bandwidth 100** is configured under the Frame-Relay interface. Only 75 percent of available bandwidth is used otherwise by default. If you have configured this correctly, as shown in Example 2-35, you have scored 3 points.

EXAMPLE 2-35 R2 QoS Configuration and Verification

```
R2(config)# class-map match-all VIDEO
R2(config-cmap)# match dscp af41
R2(config-cmap)# policy-map VIDEO-QOS
R2(config-pmap)# class VIDEO
R2(config-pmap-c)# priority percent 40
R2(config-pmap-c)# compress header ip rtp
R2(config-pmap-c)# class class-default
```

```
R2(config-pmap-c)# bandwidth percent 60
R2(config-pmap-c)# random-detect
R2(config-pmap-c)# exit
R2(config)# interface Serial0/0
R2(config-if)# max-reserved-bandwidth 100
R2(config-if)# service-policy output VIDEO-QOS
```

Section 6: Multicast (7 Points)

- Configure Routers R1, R2, R3, and R4 for IPv4 multicast. Each router should use PIM sparse dense mode. Both R1 and R2 should be configured to be Candidate RPs specifically for the following multicast groups: 225.225.0.1, 225.225.0.2, 225.225.0.3, and 225.225.0.4 by use of their Loopback 0 interfaces. You should limit the boundary of your multicast network so it does not propagate further into your network than R4. R3 should be configured as a mapping agent to announce the rendezvous points for the multicast network with the same boundary constraints. (3 points)

The question dictates that R1 and R2 be rendezvous points and advertise the same groups to the multicast network. R3 is required to announce the rendezvous points, and R4 will by default elect R2 as the RP for each group because it has the higher Loopback address compared to R1 for the same groups. TTL scoping is used within the configuration to limit the boundary of advertisements on both the candidate RPs and the discovery agent up to R4. Example 2-36 shows the required configuration and RP mappings as received on R4. If you have configured this correctly, as shown in Example 2-36, you have scored 3 points.

EXAMPLE 2-36 R1, R2, R3 and R4 Multicast Configuration and Verification

```
R1(config)# ip multicast-routing
R1(config)# interface Loopback0
R1(config-if)# ip pim sparse-dense-mode
R1(config-if)# interface Serial0/0/0
R1(config-if)# ip pim sparse-dense-mode
R1(config-if)# ip pim send-rp-announce Loopback0 scope 3 group-list GROUPS
R1(config)# ip access-list standard GROUPS
R1(config-std-nacl)# permit 225.225.0.1
R1(config-std-nacl)# permit 225.225.0.2
R1(config-std-nacl)# permit 225.225.0.3
R1(config-std-nacl)# permit 225.225.0.4
```

```
R2(config)# ip multicast-routing
R2(config)# interface Loopback0
R2(config-if)# ip pim sparse-dense-mode
R2(config-if)# interface Serial0/0
R2(config-if)# ip pim sparse-dense-mode
R2(config-if)# ip pim send-rp-announce Loopback0 scope 3 group-list GROUPS
R2(config)# ip access-list standard GROUPS
R2(config-std-nacl)# permit 225.225.0.1
R2(config-std-nacl)# permit 225.225.0.2
R2(config-std-nacl)# permit 225.225.0.3
R2(config-std-nacl)# permit 225.225.0.4
```

```
R3(config)# ip multicast-routing
R3(config)# interface Loopback0
R3(config-if)# ip pim sparse-dense-mode
R3(config)# interface GigabitEthernet0/0
R3(config-if)# ip pim sparse-dense-mode
R3(config-if)# interface Serial0/0/0
R3(config-if)# ip pim sparse-dense-mode
R3(config-if)# exit
R3(config)# ip pim send-rp-discovery lo0 scope 2
R4(config-if)# ip multicast-routing
R4(config-if)# interface GigabitEthernet0/0
R4(config-if)# ip pim sparse-dense-mode
```

```
R4# show ip pim rp mapping
PIM Group-to-RP Mappings
```

```
Group(s) 225.225.0.1/32
  RP 120.100.2.1 (?), v2v1
    Info source: 120.100.34.3 (?), elected via Auto-RP
    Uptime: 00:00:03, expires: 00:02:52
Group(s) 225.225.0.2/32
  RP 120.100.2.1 (?), v2v1
    Info source: 120.100.34.3 (?), elected via Auto-RP
    Uptime: 00:00:03, expires: 00:02:56
Group(s) 225.225.0.3/32
  RP 120.100.2.1 (?), v2v1
```

```

Info source: 120.100.34.3 (?), elected via Auto-RP
Uptime: 00:00:03, expires: 00:02:55
Group(s) 225.225.0.4/32
RP 120.100.2.1 (?), v2v1
Info source: 120.100.34.3 (?), elected via Auto-RP
Uptime: 00:00:03, expires: 00:02:55

```

- Configure R3 to ensure that R4 has a candidate RP as R1 for groups 225.225.0.1 and 225.225.0.2 and R2 for groups 225.225.0.3 and 225.225.0.4. (2 points)

As detailed in the previous example, R2 will by default become the candidate RP as selected by the discovery agent (R3) because of having a higher Loopback IP address as used in the PIM announcements compared to R1. By configuring a group-list on the discovery agent, RP announcements can be filtered. Configuring two filter lists with each candidate RP associated to them allows the discovery agent to announce two different RPs. Example 2-37 shows the required configuration, a debug of the auto-rp announcements on R3 to detail the filtering and the resulting RP mappings on R4. If you have configured this correctly, as shown in Example 2-37, you have scored 2 points.

EXAMPLE 2-37 R2 QoS Configuration and Verification

```

R3(config)# ip pim rp-announce-filter rp-list R1 group-list R1-GROUPS
R3(config)# ip pim rp-announce-filter rp-list R2 group-list R2-GROUPS
R3(config)# ip access-list standard R1
R3(config-std-nacl)# permit 120.100.1.1
R3(config-std-nacl)# exit
R3(config)# ip access-list standard R2
R3(config-std-nacl)# permit 120.100.2.1
R3(config-std-nacl)# exit
R3(config)# ip access-list standard R1-GROUPS
R3(config-std-nacl)# permit 225.225.0.1
R3(config-std-nacl)# permit 225.225.0.2
R3(config-std-nacl)# exit
R3(config)# ip access-list standard R2-GROUPS
R3(config-std-nacl)# permit 225.225.0.3
R3(config-std-nacl)# permit 225.225.0.4

R3# debug ip pim auto-rp
PIM Auto-RP debugging is
on
Auto-RP(0): Received RP-announce, from 120.100.1.1, RP_cnt 1, ht
181

```

```

Auto-RP(0): Update (225.225.0.1/32, RP:120.100.1.1), PIMv2 v1
Auto-RP(0): Update (225.225.0.2/32, RP:120.100.1.1), PIMv2 v1
Auto-RP(0): Filtered 225.225.0.3/32 for RP 120.100.1.1
Auto-RP(0): Filtered 225.225.0.4/32 for RP 120.100.1.1
Auto-RP(0): Received RP-announce, from 120.100.1.1, RP_cnt 1, ht 181
Auto-RP(0): Update (225.225.0.1/32, RP:120.100.1.1), PIMv2 v1
Auto-RP(0): Update (225.225.0.2/32, RP:120.100.1.1), PIMv2 v1
Auto-RP(0): Filtered 225.225.0.3/32 for RP 120.100.1.1
Auto-RP(0): Filtered 225.225.0.4/32 for RP 120.100.1.1

```

R4# **show ip pim rp mapping**

PIM Group-to-RP Mappings

```

Group(s) 225.225.0.1/32
  RP 120.100.1.1 (?), v2v1
    Info source: 120.100.34.3 (?), elected via Auto-RP
    Uptime: 00:00:08, expires: 00:02:52
Group(s) 225.225.0.2/32
  RP 120.100.1.1 (?), v2v1
    Info source: 120.100.34.3 (?), elected via Auto-RP
    Uptime: 00:00:08, expires: 00:02:51
Group(s) 225.225.0.3/32
  RP 120.100.2.1 (?), v2v1
    Info source: 120.100.34.3 (?), elected via Auto-RP
    Uptime: 00:00:47, expires: 00:02:12
Group(s) 225.225.0.4/32
  RP 120.100.2.1 (?), v2v1
    Info source: 120.100.34.3 (?), elected via Auto-RP
    Uptime: 00:00:47, expires: 00:02:09

```

- Configure R1 to monitor traffic forwarded through itself for traffic destined to the multicast group of 225.225.0.1. If no packet for this group is received within a single 10-second interval, ensure an SNMP trap is sent to an SNMP management station on 120.100.100.100 using a community string of “public.” (2 points)

The IP multicast heartbeat feature facilitates the monitoring of the delivery of IP multicast packets and failure notification based on configurable parameters. By configuring R1 to enable the heartbeat monitoring for the group 225.255.0.1 with the subparameters of 1 and 10, the router monitors a packet lost within 1 interval of 10 seconds and will send an SNMP trap to the SNMP host 120.100.100.100, which is required to be configured within the basic SNMP trap configuration. Example 2-38 details the required multicast heartbeat configuration and verification of the SNMP trap by issue of

a ping to 225.225.0.1 from R3. Even though R1 does not have a valid IGMP join-group for this group, traffic is still directed to it, and the heartbeat process is activated. If you have configured this correctly, as shown in Example 2-38, you have scored 2 points.

EXAMPLE 2-38 R1 Multicast Heartbeat Configuration

```
R1(config)# snmp-server host 120.100.100.100 traps public
R1(config)# snmp-server enable traps ipmulticast
R1(config)# ip multicast heartbeat 225.225.0.1 1 1 10

R1# debug snmp packets

R3# ping 225.225.0.1

R1# SNMP: Queuing packet to 120.100.100.100
SNMP: V1 Trap, ent ciscoExperiment.2.3.1, addr 120.100.100.1, gentrap 6, spectrap 1
ciscoIpMRouteHeartBeatEntry.4.225.225.0.1 = 1
ciscoIpMRouteHeartBeatEntry.5.225.225.0.1 = 0
```

Section 7: Security (7 Points)

- Allow Router R6 to passively watch the SYN connections that flow to only VLAN63 for servers that might reside on this subnet. To prevent a potential DoS attack from a flood of SYN requests, the router should be configured to randomly drop SYN packets from any source to this VLAN that have not been correctly established within 20 seconds. (2 points)

The question requires that the TCP intercept feature be configured on R6. This protects TCP servers from TCP SYN-flooding attacks with a wave of half-opened connections overwhelming the servers CPU, the result of which can effectively cause a DoS attack. The default behavior of the feature is to intercept the SYN connections to a server and effectively proxy the connection until it has been correctly established. Because you are requested to passively monitor the connection, you are required to configure the feature into watch mode by use of the global **ip tcp intercept mode watch** command. You are also requested to ensure that the feature is enabled only on VLAN 63 from any source, so an access-list is required to which the intercept features restricts its monitoring. The default behavior of the feature is to drop SYN connections based on the oldest first, but the question dictated that random connections must be dropped. This is achieved with the global command **ip tcp intercept drop-mode random**. To ensure the 20-second limit is met as opposed

to the default 30 second, adjustment of the timers is required with the global command **ip tcp intercept watch-timeout 20**. If you have configured this correctly, as shown in Example 2-39, you have scored 2 points. Use of the **show tcp intercept connections** command would be useful to verify your configuration.

EXAMPLE 2-39 R6 TCP Intercept Configuration

```
R6(config)# ip tcp intercept list 100
R6(config)# access-list 100 permit tcp any 120.100.63.0 0.0.0.255
R6(config)# ip tcp intercept mode watch
R6(config)# ip tcp intercept drop-mode random
R6(config)# ip tcp intercept watch-timeout 20
```

- Configure an ACL on R1 to allow TCP sessions generated on this router and through its Ethernet interface and to block TCP sessions from entering on its Frame-Relay interface that were not initiated on it or through it originally. Do not use the established feature within standard ACLs to achieve this, and only apply ACLs on the Frame-Relay interface. The ACL should timeout after 100 seconds of locally initiated TCP inactivity; it should also enable ICMP traffic inbound for testing purposes.(3 points)

The question requires that a reflexive ACL be configured on R1. This enables TCP traffic for sessions originating from within the network but denies TCP traffic for sessions originating from outside the network. The reflexive ACL contains only temporary entries, which are automatically created when a new TCP session is initiated. The entries are simply removed 300 seconds after the session ends by default. However, the question requires this to be modified to 100 seconds. To facilitate the reflexive ACL, you must configure a standard ACL inbound on the Frame-Relay interface, which permits the required traffic inbound to R1 and only returns traffic matching the reflexive ACL. Required traffic is of course EIGRP, PIM, IPv6 tunneling, and as directed ICMP for testing. It's a cruel question because if you forget to permit any of the required traffic inbound, you'll lose points from a previous section that you might have otherwise achieved full marks in. If you didn't know what protocol IPv6 uses, you can simply use the log option on your inbound ACL on a final deny statement. This would show you that the tunneling from R3 inbound to R1 uses IP protocol 41, which must be included in your inbound ACL.

Example 2-40 shows the required configuration and verification of the reflexive ACL. Because traffic is only evaluated by the ACL as it passes through the router, Switch 1 has been configured to belong to VLAN100 to telnet through R1 to R3 in the example. When initiated by Switch 1, the telnet session passes through the ACL **FILTER-OUT** on R1 and creates an entry in the reflexive ACL **DYNAMIC-TCP**. Real-time details can be seen by issuing the **show access-lists** command on R1. The reflexive ACL permits return traffic to the telnet session inbound on the Frame-Relay interface for the

configured inactivity interval of 100 seconds. If you have configured this correctly, as shown in Example 2-40, you have scored 4 points.

EXAMPLE 2-40 R1 Reflexive ACL Configuration and Verification

```
R1(config-if)# ip access-list extended FILTER-IN
R1(config-ext-nacl)# permit icmp any any
R1(config-ext-nacl)# permit eigrp any any
R1(config-ext-nacl)# permit pim any any
R1(config-ext-nacl)# permit tcp host 120.100.3.1 host 120.100.1.1 eq bgp
R1(config-ext-nacl)# permit 41 host 120.100.123.3 host 120.100.123.1
R1(config-ext-nacl)# evaluate DYNAMIC-TCP R1(config-
ext-nacl)# ip access-list extended FILTER-OUT R1(config-ext-
nacl)# permit tcp any any reflect DYNAMIC-TCP R1(config-ext-
nacl)# exit
R1(config)# ip reflexive-list timeout 100
R1(config)# interface Serial0/0/0
R1(config-if)# ip access-group FILTER-IN in
R1(config-if)# ip access-group FILTER-OUT out
```

```
SW1(config)# interface vlan 100
SW1(config-if)# ip add 120.100.100.100 255.255.255.0
SW1(config-if)# exit
SW1(config)# ip route 120.100.3.1 255.255.255.255 120.100.100.1
SW1(config)# exit
SW1# trace 120.100.3.1
```

```
Type escape sequence to abort.
Tracing the route to 120.100.3.1
```

```
 1 120.100.100.1 0 msec 4 msec 0 msec
 2 120.100.100.1 !A * !A
SW1# telnet 120.100.3.1
Trying 120.100.3.1 ... Open
```

```
User Access Verification
```

```
Password:
```

NOTE

The Reflexive ACL is valid only for traffic flowing through the router; as such, you might experience connectivity issues if you initiate a telnet session from R1 without manipulating the telnet source option. This behavior has no bearing on points scored and should be considered a by-product of the solution. If you face a similar question in the actual exam and telnet connectivity was required from the router you are configuring, you would specifically be instructed to ensure the correct operation of telnet on that router.

```
R3>enable
Password:
R3#

R1# show access-lists
Standard IP access list 1
 10 permit 120.100.1.0 (3 matches)
 20 permit 120.100.100.0 (3 matches)
Standard IP access list GROUPS
 10 permit 225.225.0.1
 20 permit 225.225.0.2
 30 permit 225.225.0.3
 40 permit 225.225.0.4
Reflexive IP access list DYNAMIC-TCP
 permit tcp host 120.100.3.1 eq telnet host 120.100.100.100 eq 11034 (34 matches) (time left 90)
Extended IP access list FILTER-IN
 5 permit icmp any any (150 matches)
 10 permit eigrp any any (1710 matches)
 20 permit pim any any (92 matches)
 25 permit tcp host 120.100.3.1 host 120.100.1.1 eq bgp (126 matches)
 30 evaluate DYNAMIC-TCP
Extended IP access list FILTER-OUT
 10 permit tcp any any reflect DYNAMIC-TCP (18 matches)
```

- Configure R1 so it is capable of performing SCP. The router should belong to a domain of `toughtest.co.uk`; use local authentication with a username and password of `cisco`, a key size of 768 bits, and an SSH timeout of 2 minutes and retry value of 2. (2 points).

SCP is Secure Copy Protocol; it's similar to remote copy but requires SSH to be running on the router for security purposes. It's a tough question because this is the kind of feature for which you will need to check the documentation. You will need to realize aspects of SSH are considered prerequisites to enable SCP. Even if you hadn't configured SSH or SCP previously, you should realize that you would need to configure a domain ID, local authentication with a username and password, a key of some form, and some SSH timeout and retry values based on the directions. Be careful on the values because the timeout is entered in seconds and not minutes. Your username and password combination requires a privilege level of 15 set for SCP. If you have configured this correctly, as shown in Example 2-41, you have scored 2 points.

EXAMPLE 2-41 R1 RCP Configuration

```
R1(config)# ip domain-name toughest.co.uk
R1(config)# crypto key generate rsa modulus 768
The name for the keys will be: R1.toughest.co.uk

% The key modulus size is 768 bits
% Generating 768 bit RSA keys, keys will be non-exportable...[OK]

R1(config)# aaa new-model
R1(config)# aaa authentication login default local
R1(config)# aaa authorization exec default local
R1(config)# username cisco privilege 15 password 0 cisco
R1(config)# ip ssh time-out 120
R1(config)# ip ssh authentication-retries 2
R1(config)# ip scp server enable
R1(config)#
00:57:29.343: %SSH-5-ENABLED: SSH 1.99 has been enabled
```

Lab WRAP-UP

So how did it go? Did you run out of time? Did you manage to finish but miss what was actually required? If you scored more than 80, well done. If you accomplished this within the time frame of 8 hours or less, you will be prepared for any scenario that you are likely to face during the 5 1/2 hours of the Configuration section of the actual exam. Remember that the Troubleshooting section on the v4.0 exam is a separate section to the configuration with a different scenario, and you will have 2 hours to complete this. This lab was designed to ensure you troubleshoot your own work as you progress through the questions.

Did you manage to configure items such as EIGRP third-party next hop and the continue statement within your BGP prepending? Items such as these might seem inconsequential, but they can make or break your lab.

Practice Lab 3—The VPN Lab

The CCIE exam commences with 2 hours of troubleshooting followed by 5 1/2 hours of configuration and a final 30 minutes of additional questions. This lab has been timed to last for 8 hours of configuration and self-troubleshooting, so aim to complete the lab within this period.

Then either score yourself at this point or continue until you feel you have met all the objectives. You now are going to be guided through the equipment requirements and pre-lab tasks in preparation for taking this practice lab.

If you don't own six routers and four switches, consider using the equipment available and additional lab exercises and training facilities that can be found within the CCIE R&S 360 program. Detailed information on the 360 program and CCIE R&S exam can be found on the following URLs, respectively:

https://learningnetwork.cisco.com/community/learning_center/cisco_360/360-rs

https://learningnetwork.cisco.com/community/certifications/ccie_routing_switching

NOTE

The 3825s used in this lab were loaded with c3825-adventerprisek9-mz.124-6.T.bin, and the 3725 was loaded with c3725-adventerprisek9-mz.124-6.T.bin.

NOTE

The 3550 in this lab was loaded with c3550-ipservicesk9-mz.122-25.SEE.bin, and the 3560s with c3560-ipservicesk9-mz.122-25.SEE.bin.

Equipment List

You need the following hardware and software components to begin this practice lab:

- Six routers loaded with Cisco IOS Software Release 12.4 Advanced Enterprise image and the minimum interface configuration, as documented in Table 3-1

TABLE 3-1 Hardware Required per Router

Router	Model	Ethernet I/F	Serial I/F
R1	3825	1	1
R2	3725	—	2
R3	3825	—	2
R4	3825	1	1
R5	3825	1	1
R6	3825	2	—

NOTE

Notice in the initial configurations supplied that some interfaces do not have IP addresses pre-configured. This is because either you do not use that interface or you need to configure this interface from default within the exercise. The initial configurations supplied should be used to preconfigure your routers and switches before the lab starts.

If your routers have different interface speeds than those used within this book, adjust the bandwidth statements on the relevant interfaces to keep all interface speeds in line. This ensures that you do not get unwanted behavior because of differing Interior Gateway Protocol (IGP) metrics.

- One 3550 switch with Cisco IOS Software Release 12.2 IP Services and three 3560 switches with Cisco IOS Software Release 12.2 IP Services.

Setting Up the Lab

You can use any combination of routers as long as you fulfill the requirements within the topology diagram, as shown in Figure 3-1. However, it is recommended that you use the same model of routers because this makes life easier if you load configurations directly from those supplied into your own devices.

Lab Topology

This practice lab uses the topology as outlined in Figure 3-1, which you must re-create with your own equipment.

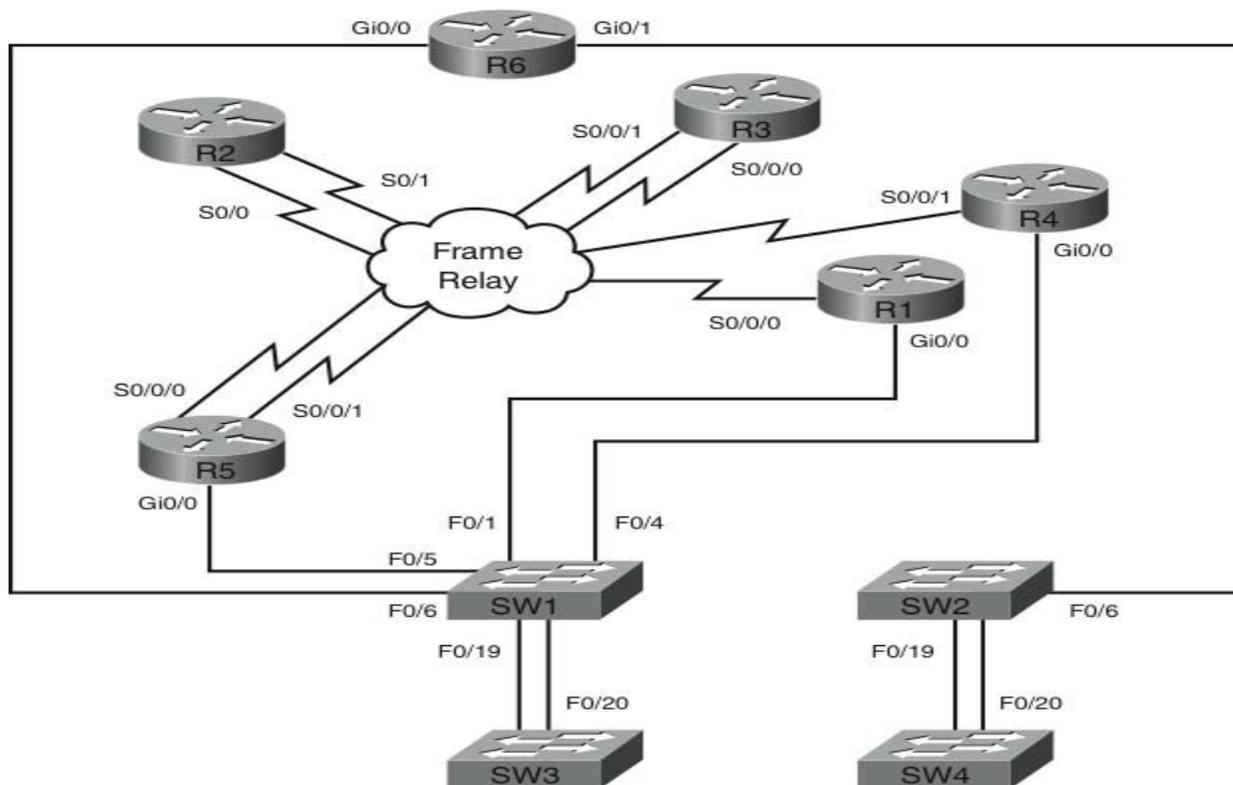


FIGURE 3-1
Lab Topology Diagram

Switch Instructions

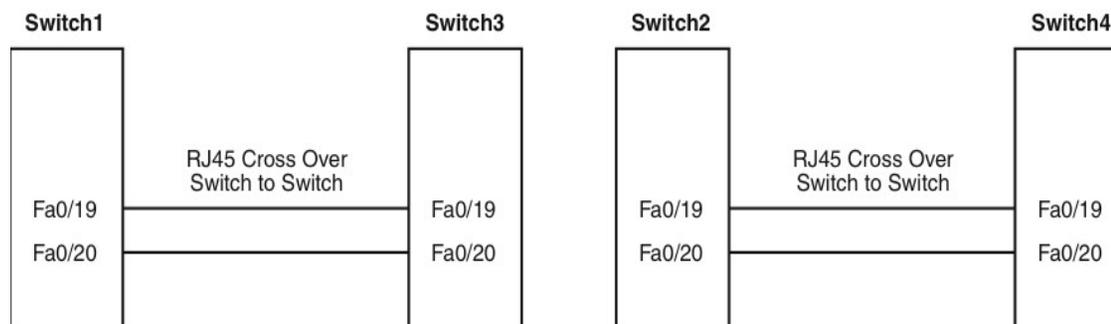
Configure VLAN assignments from the configurations supplied on the CD-ROM or from Table 3-2.

TABLE 3-2 VLAN Assignment

VLAN	Switch1	Switch2	Switch3	Switch4
45	Fa0/4, Fa0/5, Fa0/6	—	—	—
200	Fa0/19	—	—	—
400	—	Fa0/19	—	—
Trunk	Fa0/1	Fa0/6	—	—
Trunk	Fa0/20	Fa0/20	Fa0/20	Fa0/20

Connect your switches with RJ45 Ethernet Cross Over cables, as shown in Figure 3-2.

FIGURE 3-2
Switch-to-Switch Connectivity



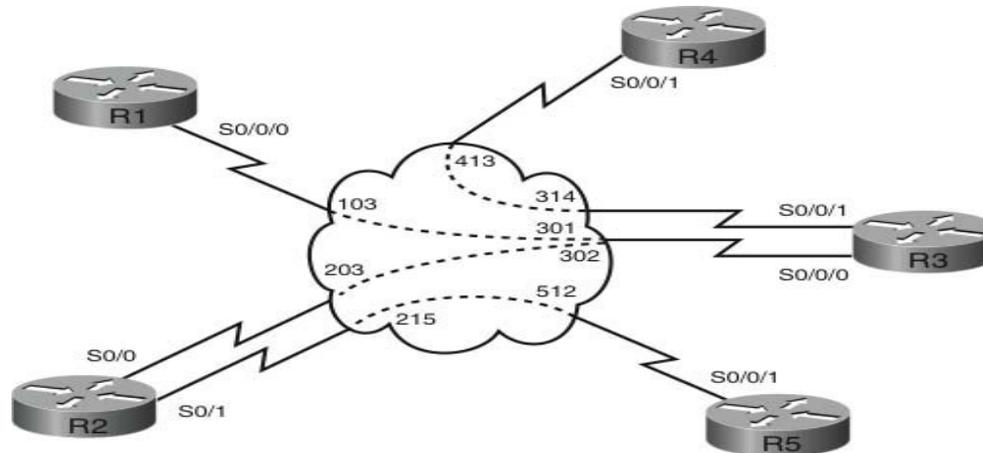
Frame Relay Instructions

Configure one of the routers you are going to use in the lab as a Frame Relay switch, or have a dedicated router purely for this task. This lab uses a dedicated router for the Frame Relay switch. A fully meshed environment is configured between all the Frame Relay routers. Pay attention in the lab as to which permanent virtual circuits (PVC) are actually required. Keep the encapsulation and Local Management Interface (LMI) settings to default for this exercise, but experiment with the settings outside these labs because you could be required to configure the Frame Relay switching within your actual lab.

If you are using your own equipment, keep the data circuit-terminating equipment (DCE) cables at the frame switch end for simplicity and provide a clock rate to all links from this end.

After configuration, the Frame Relay connectivity represents the logical Frame Relay network, as shown in Figure 3-3.

FIGURE 3-3
Frame Relay Logical
Connectivity



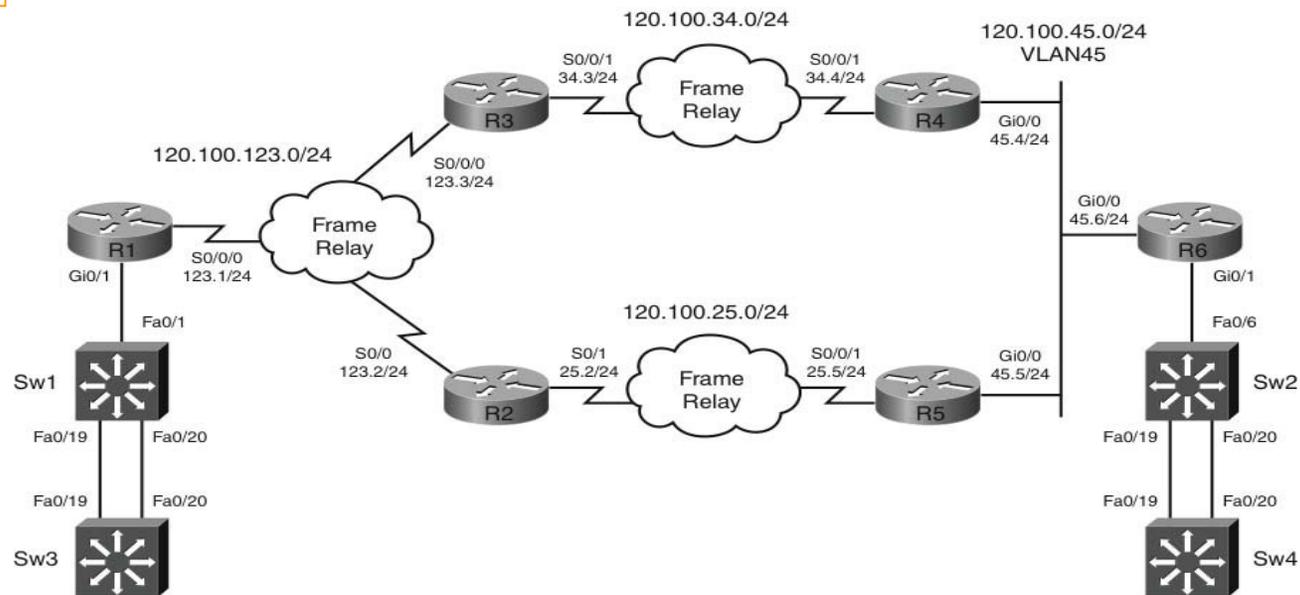
IP Address Instructions

In the actual CCIE lab, you find that the majority of your IP addresses are preconfigured. For this exercise you are required to configure your IP addresses as shown in Figure 3-4 or to load the initial router configurations supplied. If you are manually configuring your equipment, be sure you include the following loopback addresses:

```
R1 Lo0 120.100.1.1/32
R2 Lo0 120.100.2.1/32
R3 Lo0 120.100.3.1/32
R4 Lo0 120.100.4.1/32
R5 Lo0 120.100.5.1/32
R6 Lo0 120.100.6.1/32
SW1 Lo0 10.1.1.1/24
    Lo1 10.1.2.1/24
    Lo2 10.1.3.1/24
SW2 Lo0 10.2.2.1/24
```

Lo1 10.2.3.1/24
 Lo2 10.2.4.1/24
 SW3 Lo0 10.33.33.1/24
 Lo1 10.33.34.1/24
 Lo2 10.33.35.1/24
 SW4 Lo0 10.44.44.1/24
 Lo1 10.44.45.1/24
 Lo2 10.44.46.1/24

FIGURE 3-4
 IP Addressing Diagram



Pre-Lab Tasks

- Build the lab topology per Figure 3-1 and Figure 3-2.
- Configure your Frame Relay switch router to provide the necessary data-link connection identifiers (DLCI) per Figure 3-3.

- Configure the IP addresses on each router as shown in Figure 3-4 and add the loopback addresses. Alternatively, you can load the initial configuration files supplied if your router is compatible with those used to create this exercise.

General Guidelines

- Read the whole lab before you start.
- Do not configure any static/default routes unless otherwise specified.
- Use only the DLCIs provided in the appropriate figures.
- Ensure full IP visibility between routers for ping testing/Telnet access to your devices.
- If you are running out of time, choose questions that you are confident you can answer. Failing this, choose questions with a higher point rating to maximize your potential score.
- Get into a comfortable and quiet environment where you can focus for the next 8 hours.
- Take a 30-minute break midway through the exercise.
- Have available a Cisco Documentation CD-ROM, or access online the latest documentation from the following URLs: www.cisco.com/univercd/home/home.htm.

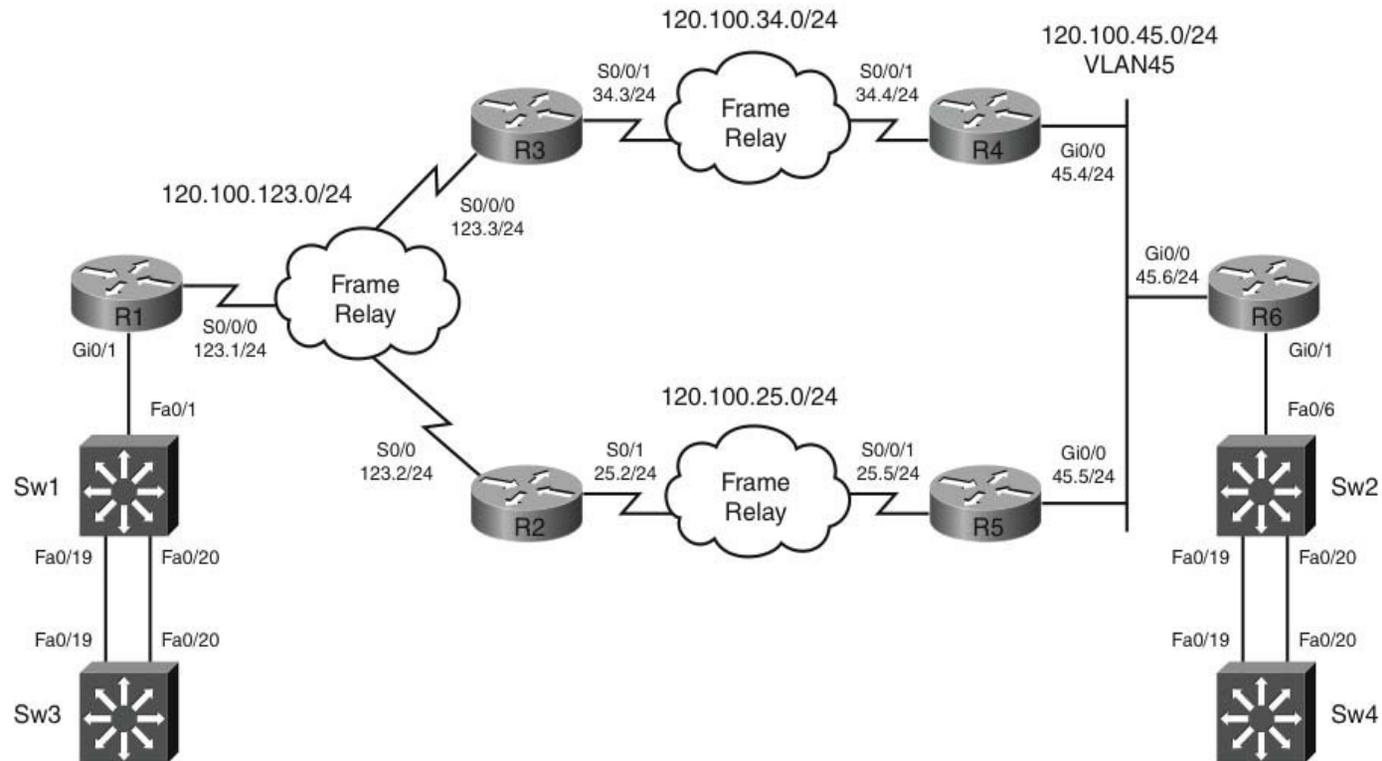
www.cisco.com/en/US/products/ps6350/products_installation_and_configuration_guides_list.html

NOTE

Access only these URLs, not the whole Cisco.com website because if you are permitted to use documentation during your CCIE lab exam, it will be restricted. Consider opening several windows with the pages you are likely to look at to save time during your lab.

Practice Lab Three

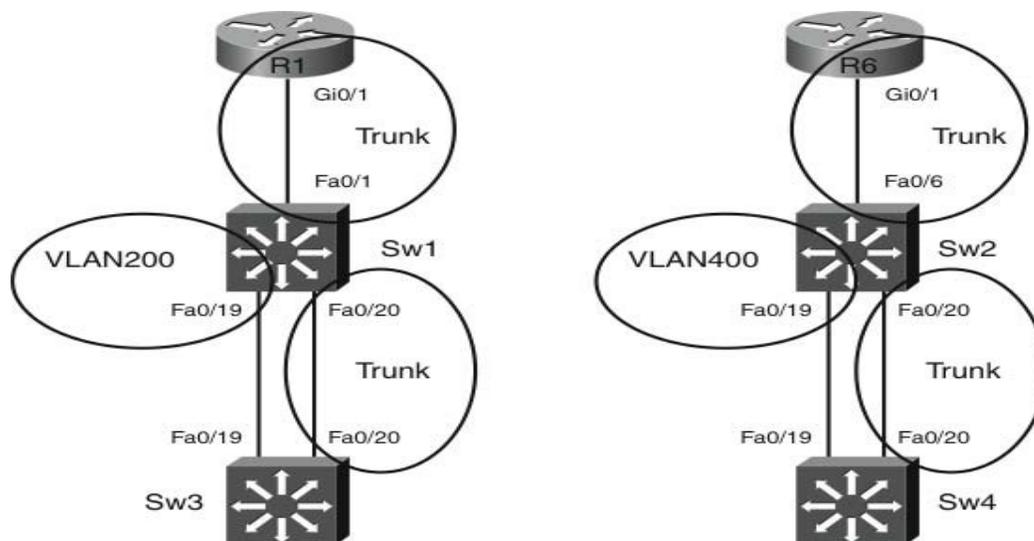
FIGURE 3-5
Lab Topology Diagram



You will now be answering questions in relation to the network topology as shown in Figure 3-5.

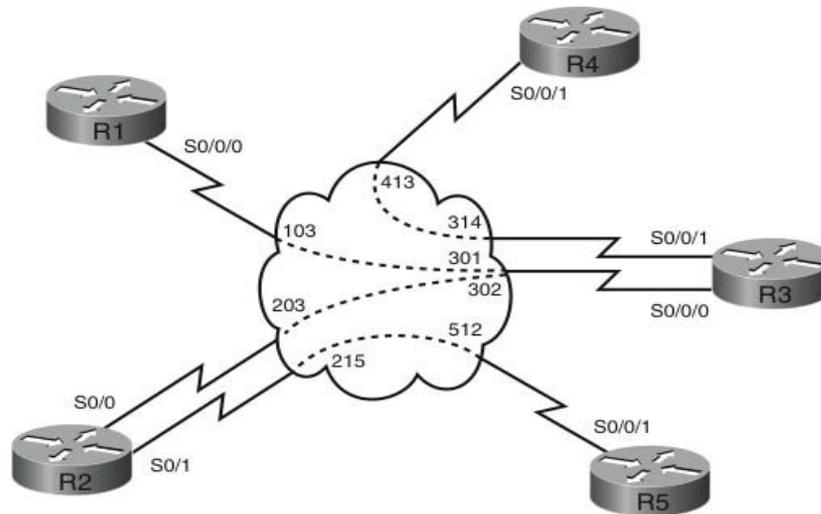
Section 1: LAN Switching and Frame Relay (6 Points)

FIGURE 3-6
Switch Topology
Diagram



- Configure your switched network per Figure 3-6. Your switched network is physically nonlooped and therefore does not require any STP root bridge configuration. Configure SW1 Fa0/19 to belong to VLAN200 and SW2 Fa0/19 to belong to VLAN400. Configure Interface Fa0/1 on SW1 to become a trunk port toward R1 and Fa0/6 on SW2 to become a trunk port toward R6; ports should use 802.1Q encapsulation. Restrict the VLANs permissible to use the trunk on Switch 1 Fa0/1 to VLAN10, 50, and 200 and VLAN20, 100 and 400 on Switch 2 Fa0/6. Interface Fa0/20 of each switch has been preconfigured to be a trunk port. You should also configure R1 and R6 to terminate the VLANs on each router. Connectivity between switches will be provided via R1 and R6 later in the lab. (3 points)
- SW3 interface Fa0/19 and SW4 interface Fa0/19 are required to communicate with each other on the same IP subnet of 1.1.1.0/24; configure these interfaces with IP addresses 1.1.1.1/24 and 1.1.1.2/24, respectively. The interfaces should be configured to communicate as if connected directly as a point-to-point link. (*Actual IP end-to-end connectivity will be achieved in a later section.*) (1 points)

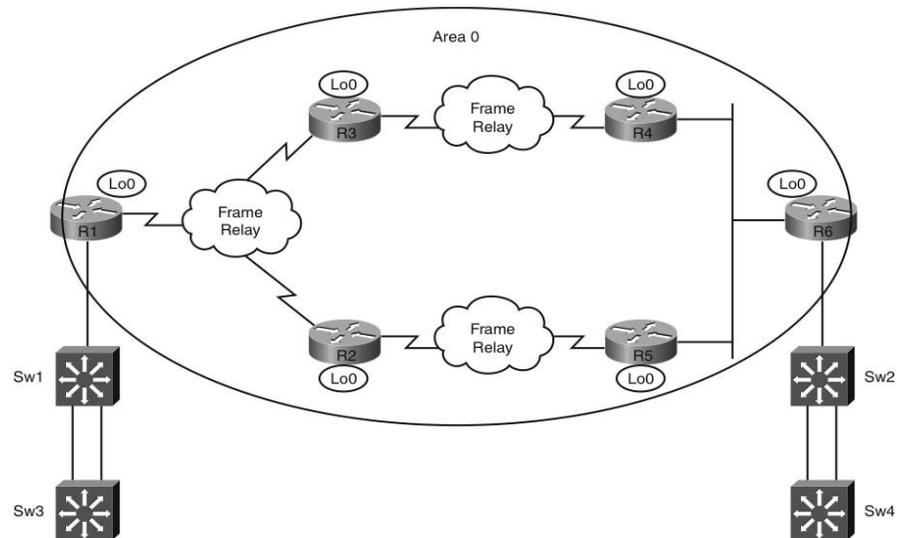
FIGURE 3-7
Frame-Relay
Connectivity Diagram



- Your initial Frame-Relay configuration has been supplied for the R1-R2-R3, R3-R4, and R2-R5 connectivity. Configure Frame-Relay per Figure 3-7 to ensure each device is reachable over the Frame-Relay network. Only use the indicated DLCIs. (2 points)

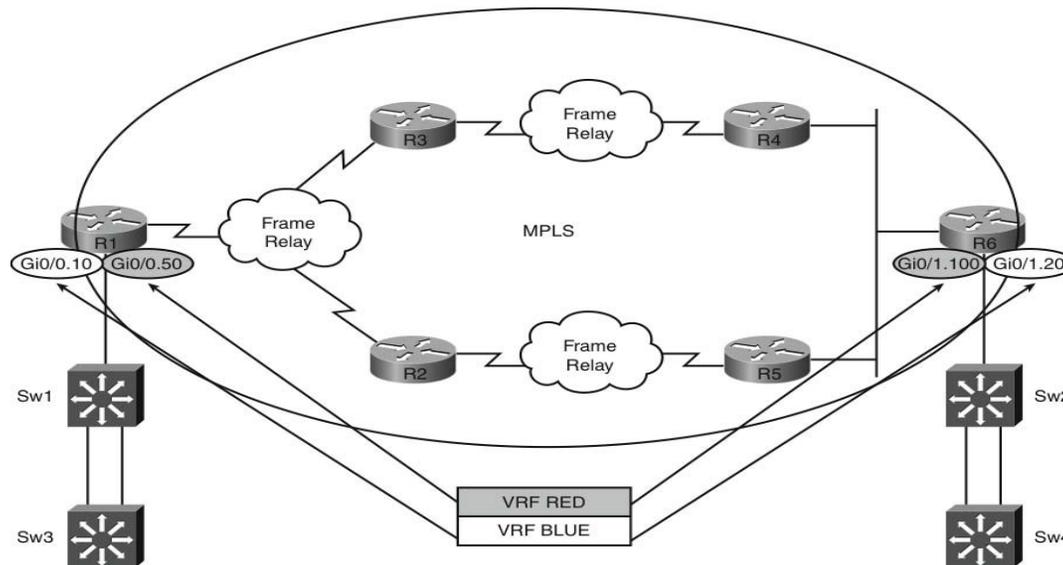
Section 2: MPLS and OSPF (19 Points)

FIGURE 3-8
Frame-Relay
Connectivity Diagram



- Configure OSPF on your routers per Figure 3-8 to enable your network to transport MPLS and MP-BGP. All required interfaces (including Loopback 0) should be configured to belong to Area 0. Ensure all OSPF configuration is entered under the interfaces. (3 points)
- Configure MPLS on all routers within the OSPF domain; use LDP, ensuring that TDP can be used on unused interfaces without specifically configuring these interfaces for TDP. Routers R1 and R6 will become your PE routers, whereas R2, R3, R4, and R5 will become P routers. (4 points)

FIGURE 3-9
VRF Topology



- You will be configuring two VPNs over your MPLS networks per Figure 3-9 between PE routers of BLUE and RED. At this point, assign the following interfaces on each PE router into separate routing instances within the routers:

PE R1 interface Gi0/0 VLAN10 connection into VPN BLUE

PE R1 interface Gi0/0 VLAN50 connection into VPN RED PE

R6 interface Gi0/1 VLAN20 connection into VPN BLUE PE

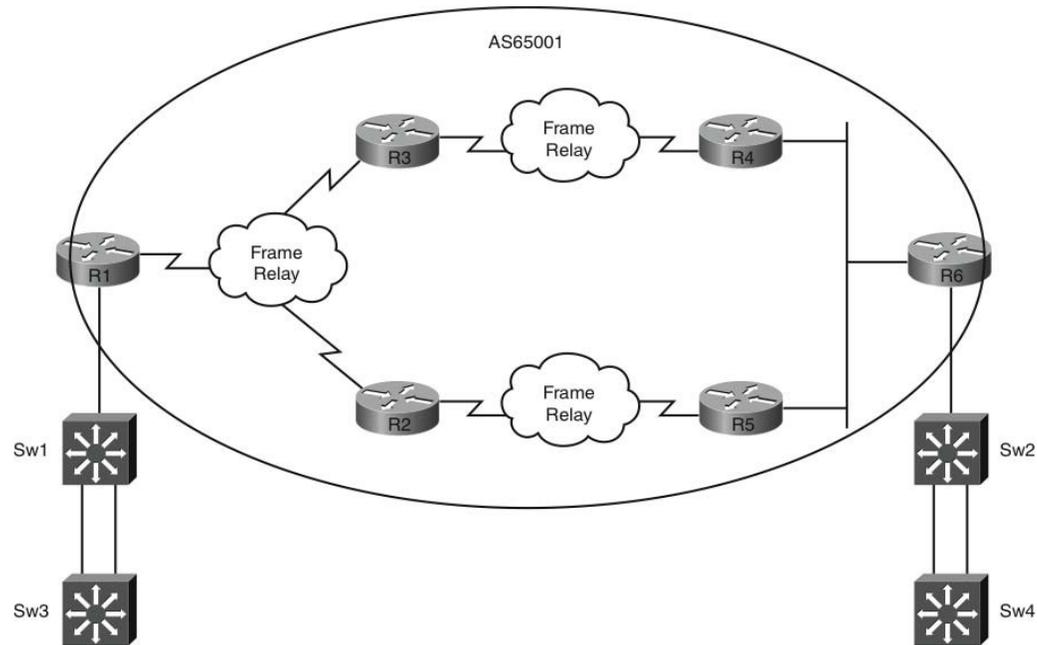
R6 interface Gi0/1 VLAN100 connection into VPN RED

Configure VPN BLUE to use an RD of 100 and VPN RED to use an RD of 200 for both importing and exporting routes into your BGP network, which will be configured later with an AS of AS65001. (4 points)

- Create a network between PE Router R1 and CE device Sw1 using a VLAN10 interface on Sw1 that can be trunked toward R1; this network will reside in the BLUE VPN. Use a subnet of 10.10.10.0/30 with .1/30 assigned to the PE and .2/30 assigned to the CE. (2 points)
- Create a network between PE Router R6 and CE device Sw2 using a VLAN20 interface on Sw2 that can be trunked toward R6; this network will reside in the BLUE VPN. Use a subnet of 10.10.20.0/30 with .1/30 assigned to the PE and .2/30 assigned to the CE. (2 points)
- Create a network between PE Router R1 and CE device Sw3 using a VLAN50 interface on Sw3 that can be trunked toward R1; this network will reside in the RED VPN. Use a subnet of 130.50.50.0/30 with .1/30 assigned to the PE and .2/30 assigned to the CE. (2 point)
- Create a network between PE Router R6 and CE device Sw4 using a VLAN100 interface on Sw4 that can be trunked toward R6; this network will reside in the RED VPN. Use a subnet of 130.100.100.0/30 with .1/30 assigned to the PE and .2/30 assigned to the CE. (2 points)

Section 3: BGP (5 Points)

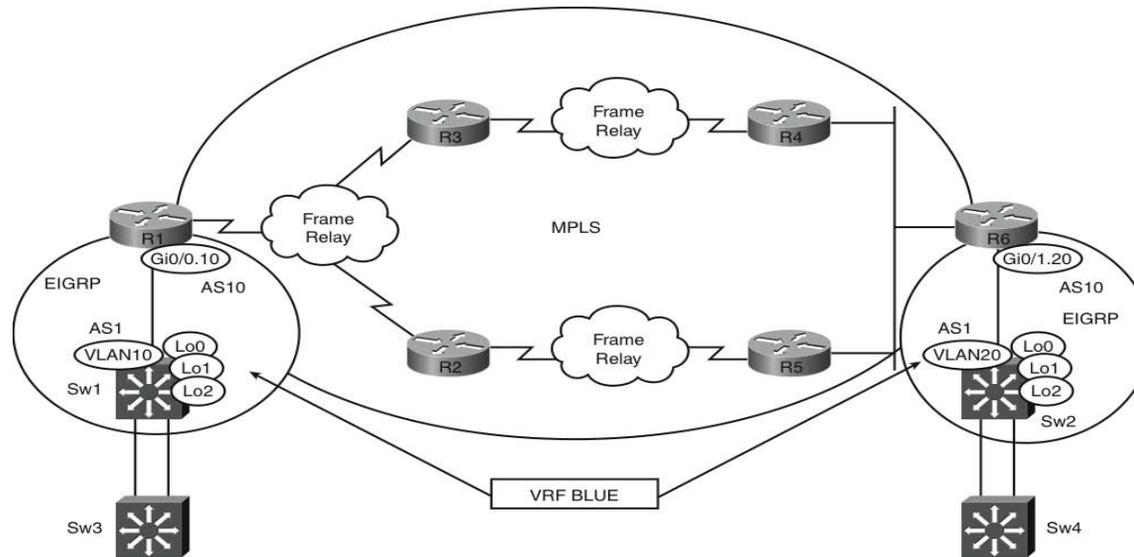
FIGURE 3-10
BGP Topology



- Configure MP-BGP between your PE routers, per Figure 3-10, to enable your network to transport the VPNv4 addresses of your configured VPNs (BLUE and RED). Use loopback interfaces for peering between your PE routers. You will configure the actual VPN routing in later questions. (4 points)

Section 4: EIGRP and MP-BGP (9 Points)

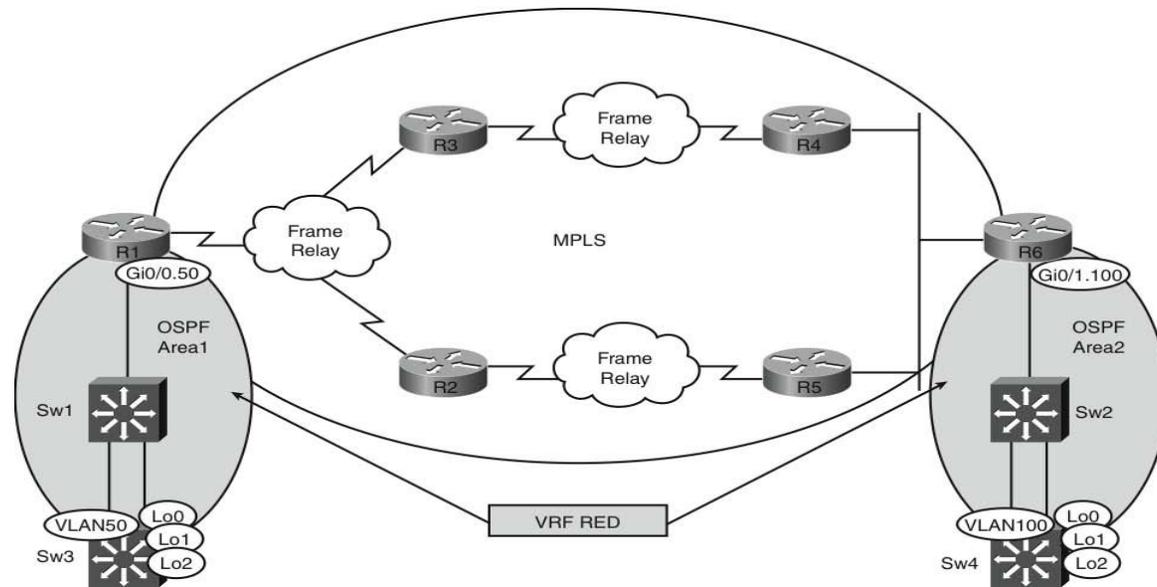
FIGURE 3-11
EIGRP Topology



- Configure EIGRP per Figure 3-11 between your PE Router R6 and CE Switch Sw2. Use an EIGRP process number of 1 on R6 and a process number of 10 on Sw2. Use VLAN20 for EIGRP connectivity between R6 and Sw2. Advertise all preconfigured Loopback networks on Sw2 to R6 for the BLUE VPN. (3 points)
- Configure EIGRP per Figure 3-11 between your PE Router R1 and CE Switch Sw1. Use an EIGRP process number of 1 on R1 and a process number of 10 on Sw1. Use VLAN10 for EIGRP connectivity between R1 and Sw1. Advertise all preconfigured Loopback networks on Sw1 to R1 for the BLUE VPN. (3 points)
- Configure your PE Routers R1 and R6 to transport EIGRP routes from your CE devices between the BLUE VPN using MP-BGP. EIGRP networks residing on Sw1 should be seen as internal EIGRP routes on Sw2 and vice versa. Ensure all EIGRP routes have a MED of 50 assigned to them within MP-BGP. Use a default-metric of 10000 100 255 1 1500 for BGP routes when redistributed into EIGRP. (3 points)

Section 5: OSPF and MP-BGP (9 Points)

FIGURE 3-12
OSPF Topology



- Configure OSPF per Figure 3-12 for your VRF RED with a process number of 3 on PE Router R1 and Sw3 using VLAN50 for connectivity. Use a process ID of 2 on PE Router R6 and CE device Sw4 using VLAN100 for connectivity. You should permit only internal OSPF routes to be advertised across your VPN and ensure the redistribution of BGP routes into OSPF are assigned as Type 1 External routes with no manually adjusted cost associated to them. It is acceptable for these routes to come through as /32 routes because of default OSPF behavior of Loopback interfaces. (3 points)
- You will notice that your OSPF IA (Intra Area) routes between CE devices Sw3 and Sw4 appear as Type 1 External routes; configure your OSPF network appropriately to ensure the routes are displayed correctly as IA routes. You are not permitted to adjust the OSPF redistribution into BGP as directed in the previous question. Maintain the OSPF process IDs as previously directed; you are permitted to configure only Router R1. (6 points)

Section 6: MPLS (7 Points)

- Leak network 10.1.1.0/24 from Sw1 VRF BLUE on PE R1 into the VRF RED on PE1; similarly, leak 10.44.44.0/24 from VRF RED into VRF BLUE on R6. Both Switch 1 and Switch 4 should receive the following routes:

```
SW1# show ip route | include 10.44.44.0
```

```
D EX 10.44.44.0/24 [170/XXXXXX] via 10.10.10.1, 00:00:27, Vlan10
```

```
SW1#
```

```
SW4# show ip route | include 10.1.1.0
```

```
O E1 10.1.1.0/24 [110/XX] via 130.100.100.1, 00:03:04, Vlan100
```

```
SW4#
```

Verify your configuration by pinging from VRF RED Sw4 10.44.44.1 to VRF BLUE Sw1 10.1.1.1 sw1. (5 points)

- Configure your PE Routers R1 and R6 to ensure that the MPLS P routers are not listed as intermediate hops when a trace route is performed on your CE devices. (2 Points)

Section 7: VPLS Simulation (10 Points)

- Switches 3 and 4 will have been configured to belong to the subnet of 1.1.1.0/24 within a previous question. Create an Xconnect attachment circuit on your PE Routers R1 and R6 for your CE devices (Sw3 Fe 0/19 1.1.1.1/24 and Sw4 Fe 0/19 1.1.1.2/24) to communicate using a secure Layer 2 tunneling solution (use version 3) across your Layer 3 network. You should use existing loopback interfaces on your PE routers for peering over your MPLS network. Use a class template that configures a cookie size of 8 and a password of cisco, which will be used by a pseudowire class that Xconnects your required interfaces on your PE Routers R1 and R6. Be aware that the Sw3 resides in VLAN200, and Sw4 resides in VLAN400 in respective PE router subinterfaces. (10 points)

Section 8: Multicast (10 Points)

- Configure your MPLS network for multicast support of the RED VRF using PIM sparse mode. PE Routers R1 and R6 should be configured to tunnel multicast traffic using an MDT address of 232.0.0.11 from CE device Switch 3 VLAN50 to CE device Sw4 VLAN100 over the RED VRF. Switch 4 should be configured to reply to

an ICMP ping on its VLAN100 interface directed to 226.2.2.2 from Switch 3 VLAN50. It can be assumed that the mVRF bandwidth requirement is low; configure MDT appropriately. Ensure that PE Router R6's associated VLAN100 IP address is used as the rendezvous point for the RED VRF multicast traffic. (10 points)

Section 9: IPv6 (6 Points)

- Configure the following IPv6 address on the PE Routers R1 and R6, and implement IPv6 over MPLS between the 6PE routers to advertise the prefixes between 6PEs. Ensure your loopback IPv6 addresses are used to source any locally generated IPv6 traffic. (6 points)

R1 Lo0 2010:C15:C0:1::1/64

R1 Gi0/0.10 2010:C15:C0:11::1/64

R6 Lo0 2010:C15:C0:6::1/64

R6 Gi1/0.20 2010:C15:C0:62::1/64

Section 10: QoS (13 Points)

- Create the following QoS profile on your PE Router R1 for traffic egressing to your CE device connected to the BLUE VRF; use an appropriate method of prioritizing DSCP traffic so that AF31 packets are statistically dropped more frequently than AF32 during congestion, and reduce the effects of TCP global synchronization within your MISSION-CRITICAL and solely reduce the effect of TCP global synchronization within the DEFAULT class: (7 points)

Class	DSCP Value	% of Bandwidth Assigned
VOICE	EF, CS5	35
MISSION-CRITICAL	CS6, AF31, AF32, CS3	40
DEFAULT	Any	25

- Create the following QoS profile on your PE Router R1 for traffic ingressing from your CE device connected to the BLUE VRF into the MPLS network; the total aggregate speed from the CE to PE should be restricted to 1 Mbps:

Class	CIR (bps)
VOICE	350,000

MISSION-CRITICAL	400,000
DEFAULT	250,000

- Traffic in the VOICE class within the detailed CIR should have the MPLS EXP set to 5 and above discarded. Traffic in the MISSION-CRITICAL class within the detailed CIR should have the MPLS EXP set to 3 and above set to 7. Traffic in the DEFAULT class within the detailed CIR should have the MPLS EXP set to 0 and above set to 4. (6 points)

Section 11: Security (13 Points)

- Create three new loopback IP addresses of loopback1 on R4, R5, and R6—use IP addresses of 4.4.4.4/24, 5.5.5.5/24, and 6.6.6.6/24, respectively. Use EIGRP to advertise the loopback networks between routers over a common GRE tunnel network of 100.100.100.X/24 (X=router number) sourced from each router's common Ethernet interface using IPsec to encrypt all traffic between the loopback networks using a preshared isakmp key of CCIE. Use an IPsec transform-set of esp-des esp-md5-hmac on each router. R6 is to be a hub router with R4 and R5 being effectively spoke routers in your solution. You are not permitted to enable EIGRP on your Ethernet interfaces between routers. Spoke routers must communicate with each other directly using dynamic IPsec connections with the aid of NHRP at the hub, whereas hub-to-spoke IPsec connections should be permanent. The hub router should provide all necessary direct next-hop information to the spoke routers when they are required to communicate between themselves. NHRP should be authenticated with a password of SECRET. Use an MTU of 1416 for your secure traffic, an NHRP timeout of 100 seconds for spoke replies, and a delay of 2mS on the tunnel network. Test your solution by extended pings sourced from the configured loopback interfaces. (10 points)
- The network manager of your network cannot justify a full security implementation but wants to implement a solution that provides a password prompt from R1 only when the keyboard entry 1 is entered on the console port (as opposed to the normal CR/Enter key). Configure R1 appropriately. 3 points

NOTE

This section should be used only if you require clues to complete the questions. In the actual CCIE lab, the proctor will not enter into any discussions about the questions or answers; he or she will be present to ensure you do not have problems with the lab environment and to maintain the timing element of the exam.

Practice Lab 3: “Ask the Proctor”

Section 1: LAN Switching and Frame Relay

Q: Do you want me to configure Layer 2 between Switch 3 and Switch 4 so that they can communicate on the subnet 1.1.1.0/24?

A: No, simply configure the switches as directed in the question and Layer 2 connectivity will be provisioned later within the lab when your core network is configured.

Q: With my Frame Relay I can only reach my spoke routers from the Hub. Is this acceptable?

A: No, the question states that each device must be reachable over the frame-relay network; this includes spoke-to-spoke communication.

Section 2: MPLS and OSPF

Q: Do you require OSPF for any interfaces on R1 and R6 that connect to the switches?

A: No, just configure OSPF per the figure; this is required to advertise your loopback addresses for MPLS.

Q: Does it matter what OSPF Process ID I use on my routers?

A: No, the question doesn't direct you to use a specific process ID, so you can use an ID of your choice.

Q: Do you want the OSPF from the core routers extended into the RED VRF I created so I run end-to-end OSPF between CE Switch1 and CE switch2?

A: No, you will ultimately achieve this connectivity through an MPLS VPN and not by simply extending OSPF through your core devices.

Q: Do you want me to configure my RED VRF with a route descriptor of 100 and 200 for the BLUE VRF?

A: You have been provided with additional information in the question that enables you to facilitate use of MP-BGP extended communities.

Q: So just add in the MP-BGP AS number to the RD?

A: A combination of the two will achieve the desired results.

Q: I can't ping to my VLAN10 interface on Switch1 from R1. Do I need to perform any further configuration to make this work?

A: No, just remember that R1 is now a PE router with multiple VRF routing tables. You need to ensure you source your ping correctly; otherwise, R1 would use its default routing table (which is used for the MPLS connectivity).

Section 3: BGP

Q: Do you want me to configure a full mesh of BGP between all routers?

A: No, MP-BGP is simply required between the PE routers.

Q: Do you need me to configure the PEs to send community values to each other?

A: You need to remember how MPLS works and ensure that the route targets are propagated to successfully configure your VPNs.

Q: I usually configure next-hop self on my BGP configurations. Is this acceptable here?

A: You haven't been instructed not to use this command at this point even though this is an iBGP configuration.

Section 4: EIGRP and MP-BGP

Q: EIGRP requires the same AS number on neighbor routers to peer successfully. If I use a different number on R6 and Switch2, they cannot peer correctly.

A: Correct. Look for a method of making the AS number the same within your VRF specific configuration on R6.

Section 5: OSPF and MP-BGP

Q: Do you want me to configure OSPF, MPLS, and BGP initially within the OSPF section?

A: No, just initially as directed OSPF; this will enable your network to transport MPLS and BGP within later questions.

Q: Changing the process ID on OSPF peers wouldn't affect any adjacency. Why would I need to do this?

A: You are correct, but you have been directed to do so in the question. It will become evident why you have been asked to do this in a later question.

Q: Why would I want to advertise the OSPF routes as External type-1 routes within BGP; surely the routes should appear as standard interarea routes through the VPN?

A: Correct, this question is a little misleading. The routes will come out as Type-1 External routes on your CE devices, and it would appear that you have modified this behavior with your redistribution configuration. This behavior should become apparent why in the following question.

Q: I think if I change the redistribution of OSPF into BGP, I can make the OSPF routes appear as Intra-area routes. Do I score any points if I change the redistribution?

A: No, by all means try to change the redistribution, though; it might help you understand the issue.

Q: I changed the redistribution and the routes remain identical. This must have something to do with the different OSPF process ID I had to configure; I can't adjust this, so I am stuck.

A: You had a similar issue with EIGRP AS numbers; just investigate what is possible within your VRF configuration.

Q: If I change the domain ID on R1, is that acceptable?

A: Find an appropriate value and try it out.

Section 6: MPLS

Q: I can manage to leak routes between VRFs but my route comes out as a host route. Can I modify my Loopback interface with the OSPF **network** command on Switch4, so it is advertised with the correct mask?

A: Yes

Section 7: VPLS Simulation

Q: Do you want me to create a pseudowire with MPLS encapsulation to connect Switch3 and Switch4 at Layer 2?

A: No, you might have found this question in the MPLS section if that were wanted; the clue is in the question as to which solution you should use.

Q: Is this MPLS-specific, or could I do this over a standard Layer 3 network?

A: You could achieve the same result over a standard Layer 3 network; just exercise caution where you configure your parameters to achieve the correct results in the appropriate VRF.

Q: Xconnect is usually associated with L2TP. Can I use this technology for my solution?

A: Yes.

Q: I have my L2TPv3 tunnel up end-to-end, yet I cannot ping between switches. I suspect a spanning-tree type issue if the question states VLAN differences when I need to provide Layer 2 adjacency. Am I at liberty to manipulate spanning tree?

A: Yes.

Section 8: Multicast

Q: Do you want me to enable PIM over my P routers or just PE routers?

A: The question states “MPLS network.” To provide end-to-end multicast support, you might find that configuring PIM end-to-end is required.

Q: Do you want PIM on my MPLS router loopback interfaces?

A: You might find it is required at certain points within your MPLS network.

Q: I have a Multicast Distribution Tree tunnel between PE routers, but I don’t understand what the low bandwidth requirement is.

A: MDT has differing requirements for high and low bandwidth sources; you might or might not require a Data MDT.

Q: To get Switch 4 to reply to a ping to 226.2.2.2, can I just configure an IGMP join group appropriately on its VLAN100 interface?

A: You can.

Section 9: IPv6

Q: Do you want me to run IPv6 down to my CE switches and redistribute anything over MPLS?

A: Your switches are currently not capable of running IPv6.

Q: Should I just advertise my IPv6 prefixes with the BGP network command?

A: Yes, because there is no redistribution to be configured.

Section 10: QoS

Q: Do you want the first QoS policy outbound on the BLUE VRF interface on PE Router R1?

A: Yes.

Q: To prioritize DSCP traffic, do you want me to configure some priority queuing within a class for AF32 flows?

A: No, use a common technique whereby traffic is dropped randomly as queues fill. AF31 packets should be dropped more frequently than AF32, though.

Q: Are you looking for Random Early Detect?

A: You're almost there; this wouldn't offer the inherent drop preference, though.

Q: The second QoS policy limits traffic to 1 Mbs, yet the first will be line rate at 1 Gbp. Is this correct?

A: Yes, I appreciate that this isn't the real world; it just provides you with two different configuration exercises.

Q: Do I use the same packet marking classes in each question?

A: Yes.

Q: Is this DiffServ whereby you want me to modify the topmost bits in the EXP field?

A: Yes.

Q: Do you want the policy applied to the CE facing VRF BLUE interface as an input service policy?

A: Yes, this would then modify the traffic as it flows into the MPLS network.

Section 11: Security

Q: Don't I need an ACL to mark all traffic that should be encrypted?

A: No, your solution will not require an ACL, and all traffic flowing from the new subnets you created should automatically be encrypted.

Q: The clues in the question suggest this is a DMVPN question. I have configured my solution correctly, yet I don't get spoke routes on the spoke routers. Is this acceptable?

A: No, you need full network visibility from all devices and not just the hub.

Q: This sounds like a split-horizon issue; can I disable this behavior?

A: Yes.

Q: I still show a next hop of the hub between spoke networks, is this okay?

A: No; the question specifically states that spoke routers must be able to communicate with each other directly.

Q: Can I modify the next hop from the hub?

A: Yes.

Q: Do you want me to get R1 to somehow translate a CR into a 1 to then provide a password prompt?

A: No, just make the router provide a prompt when it receives an ASCII 1, rather than a CR on the line con 0 port.

Practice Lab 3 Debrief

The lab debrief section now analyzes each question showing you what was required and how to achieve the desired results. You should use this section to produce an overall score for Practice Lab 3.

Section 1: LAN Switching and Frame Relay (6 Points)

- Configure your switched network per Figure 3-6. Your switched network is physically nonlooped and therefore does not require any STP root bridge configuration. Configure SW1 Fa0/19 to belong to VLAN200 and SW2 Fa0/19 to belong to VLAN400. Configure Interface Fa0/1 on SW1 to become a trunk port toward R1 and Fa0/6 on SW2 to become a trunk port toward R6; ports should use 802.1Q encapsulation. Restrict the VLANs permissible to use the trunk on Switch 1 Fa0/1 to VLAN10, 50, and 200 and VLAN20, 100, and 400 on Switch 2 Fa0/6. Interface Fa0/20 of each switch has been preconfigured to be a trunk port. You should also configure R1 and R6 to terminate the VLANs on each router. Connectivity between switches will be provided via R1 and R6 later in the lab. (3 points)

This is a simple question, but you are required to complete multiple configuration items to gain your points. The configuration enables connectivity between switches when the MPLS section has been completed later in the lab. To begin, Ports Fa0/19 of Switch 1 and Switch 2 should be assigned the correct VLAN. (The actual VLANs would have been created previously in the initial configuration.) Next, the trunking is configured as directed with allowed VLANs of 10, 50, and 200 for Switch 1 and 20, 100, and 400 for Switch 2. R1 and R6 are configured with the corresponding VLAN numbers as subinterfaces to terminate the trunk connections from switch1 and switch2 using an identical reference for the dot1q encapsulation. If you have configured this correctly as shown in Example 3-1, you have scored 3 points.

EXAMPLE 3-1 Sw1, Sw2, R1, and R6 Configuration

NOTE

R1 and R6 use the VLAN number for the encapsulation and the sub interface number. Your sub interface number does not need to match the VLAN number, but it is considered good practice to do so.

```
Switch1# show run interface FastEthernet 0/19
!
interface FastEthernet0/19
  switchport access vlan 200
  switchport mode access

Switch1# show run interface FastEthernet 0/1
!
interface FastEthernet0/1
  switchport trunk encapsulation dot1q
  switchport trunk allowed vlan 10,50,200
```

```
switchport mode trunk
```

```
Switch2# show run interface FastEthernet 0/19
```

```
!  
interface FastEthernet0/19  
  switchport access vlan 400  
  switchport mode access
```

```
Switch2# show run interface FastEthernet 0/6
```

```
!  
interface FastEthernet0/6  
  switchport trunk encapsulation dot1q  
  switchport trunk allowed vlan 20,100,400  
  switchport mode trunk
```

```
R1# show run | begin interface GigabitEthernet0/0
```

```
!  
interface GigabitEthernet0/0  
  no ip address  
!  
interface GigabitEthernet0/0.10  
  encapsulation dot1Q 10  
!  
interface GigabitEthernet0/0.50  
  encapsulation dot1Q 50  
!  
interface GigabitEthernet0/0.200  
  encapsulation dot1Q 200
```

```
R6# show run | begin interface GigabitEthernet0/1
```

```
!  
interface GigabitEthernet0/1  
  no ip address  
!  
interface GigabitEthernet0/1.20  
  encapsulation dot1Q 20  
!  
interface GigabitEthernet0/0.100  
  encapsulation dot1Q 100
```

```
!  
interface GigabitEthernet0/1.400  
encapsulation dot1Q 400
```

- SW3 interface Fa0/19 and SW4 interface Fa0/19 are required to communicate with each other on the same IP subnet of 1.1.1.0/24. Configure these interfaces with IP addresses 1.1.1.1/24 and 1.1.1.2/24, respectively. The interfaces should be configured to communicate as if connected directly as a point-to-point link. (*Actual IP end-to-end connectivity will be achieved in a later section.*) (1 points)

A straightforward configuration task to change the operation of the ports to nonswitchport Layer 3 mode where an IP address can be configured, end-to-end connectivity is achieved through the IP network at a later stage. If you have configured this correctly, as shown in Example 3-2, you have scored 1 point.

EXAMPLE 3-2 Sw3 and Sw4 Configuration

```
Switch3# show run interface FastEthernet 0/19  
!  
interface FastEthernet 0/19  
no switchport  
  
ip address 1.1.1.1 255.255.255.0  
  
Switch4# show run interface FastEthernet0/1  
  
interface FastEthernet0/19  
  
no switchport  
  
ip address 1.1.1.2 255.255.255.0
```

- Your initial Frame-Relay configuration has been supplied for the R1-R2-R3, R3-R4, and R2-R5 connectivity. Configure Frame-Relay, per Figure 3-7, to ensure each device is reachable over the Frame-Relay network. Only use the indicated DLCIs. (2 points)

The initial Frame-Relay configuration has been supplied for you; all you need to do is create additional maps on R1 and R2 spoke routers to enable them to communicate with each other by directing traffic toward the Hub Router R3 (because the initial configuration uses no inverse arp). If you have configured this correctly, as shown in Example 3-3, you have scored 2 points.

EXAMPLE 3-3 R1 and R2 Additional Frame-Relay Configuration and Verification

```
R1(config)# interface Serial0/0/0  
R1(config-if)# frame-relay map ip 120.100.123.2 103 broadcast
```

```
R2(config)# interface Serial0/0
R2(config-if)# frame-relay map ip 120.100.123.1 203 broadcast

R1# ping 120.100.123.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 120.100.123.2, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 8/8/8 ms
```

Section 2: MPLS and OSPF (19 Points)

- Configure OSPF on your routers, per Figure 3-8, to enable your network to transport MPLS and MP-BGP. All required interfaces (including Loopback 0) should be configured to belong to Area 0. Ensure all OSPF configuration is entered under the interfaces. (3 points)

OSPF is used as the IGP in which to advertise the router loopback addresses, which will of course be used for the MPLS connectivity. The question directs you to configure OSPF directly under the interfaces of the routers. Example 3-4 shows the Loopback interfaces of each router from R1's perspective advertised as host routes as required for MPLS. If you have configured this correctly, as shown in Example 3-4, you have scored 3 points. Consider using the **show ip ospf interface** command to verify your configuration.

EXAMPLE 3-4 OSPF Configuration and Verification

```
R1(config-if)# int lo0
R1(config-if)# ip ospf 1 area 0
R1(config-if)# int s0/0/0
R1(config-if)# ip ospf 1 area 0

R2(config-if)# int lo0
R2(config-if)# ip ospf 1 area 0
R2(config-if)# int s0/0
R2(config-if)# ip ospf 1 area 0
R2(config-if)# int s0/1
R2(config-if)# ip ospf 1 area 0

R3(config-if)# int lo0
```

```
R3(config-if)# ip ospf 1 area 0
R3(config-if)# int s0/0/0
R3(config-if)# ip ospf 1 area 0
R3(config-if)# int s0/0/1
R3(config-if)# ip ospf 1 area 0
```

```
R4(config-if)# int lo0
R4(config-if)# ip ospf 1 area 0
R4(config-if)# int gi0/0
R4(config-if)# ip ospf 1 area 0
R4(config-if)# int s0/0/1
R4(config-if)# ip ospf 1 area 0
```

```
R5(config-if)# int lo0
R5(config-if)# ip ospf 1 area 0
R5(config-if)# int gi0/0
R5(config-if)# ip ospf 1 area 0
R5(config-if)# int s0/0/1
R5(config-if)# ip ospf 1 area 0
```

```
R6(config-if)# int lo0
R6(config-if)# ip ospf 1 area 0
R6(config-if)# int gi0/0
R6(config-if)# ip ospf 1 area 0
```

```
R1# show ip route ospf
    120.0.0.0/8 is variably subnetted, 12 subnets, 2 masks
O       120.100.25.0/24 [110/128] via 120.100.123.2, 00:34:18, Serial0/0/0
O       120.100.5.1/32 [110/129] via 120.100.123.2, 00:34:18, Serial0/0/0
O       120.100.4.1/32 [110/129] via 120.100.123.3, 00:34:18, Serial0/0/0
O       120.100.6.1/32 [110/130] via 120.100.123.3, 00:34:18, Serial0/0/0
        [110/130] via 120.100.123.2, 00:34:18, Serial0/0/0
O       120.100.3.1/32 [110/65] via 120.100.123.3, 00:34:18, Serial0/0/0
O       120.100.2.1/32 [110/65] via 120.100.123.2, 00:34:18, Serial0/0/0
O       120.100.45.0/24 [110/129] via 120.100.123.3, 00:34:18, Serial0/0/0
        [110/129] via 120.100.123.2, 00:34:18, Serial0/0/0
O       120.100.34.0/24 [110/128] via 120.100.123.3, 00:34:18, Serial0/0/0
```

```
O      120.100.123.3/32 [110/64] via 120.100.123.3, 00:34:18, Serial0/0/0
O      120.100.123.2/32 [110/64] via 120.100.123.2, 00:34:18, Serial0/0/0
```

- Configure MPLS on all routers within the OSPF domain; use LDP ensuring that TDP can be used on unused interfaces without specifically configuring these interfaces for TDP. Routers R1 and R6 will become your PE routers, whereas R2, R3, R4, and R5 will become P routers. (4 points)

Configuration is required on each router for them to become LSRs (Label Switch Routers). The LSRs must have Loop-back interfaces with an address mask of 32 bits, and these interfaces must be reachable within the global IP routing table (which the previous question achieved). R1 and R6 are the PE (Provider Edge) routers, which will be used to connect to switches in later questions simulating CE (Customer Edge) devices. R2, R3, R4, and R5 become the P (Provider) routers, which will be used to switch labeled packets between the PE routers. The question tells you to use LDP (Label Distribution Protocol) but facilitate the future use of TDP (Tag Distribution Protocol) without further configuration on unused interfaces. This is achieved by configuring TDP globally and LDP under each interface used for MPLS within this lab. (The default global and interface configuration is LDP). The PE routers require only MPLS configured on their serial interfaces toward the P routers. If you have configured this correctly, as shown in Example 3-5, you have scored 4 points.

EXAMPLE 3-5 MPLS Configuration

```
R1(config)# mpls label protocol tdp
R1(config)# interface Serial0/0/0
R1(config-if)# mpls label protocol ldp
R1(config-if)# mpls ip

R2(config)# mpls label protocol tdp
R2(config)# interface Serial0/0
R2(config-if)# mpls label protocol ldp
R2(config-if)# mpls ip
R2(config-if)# mpls label protocol ldp
R2(config-if)# mpls ip

R3(config)# mpls label protocol tdp
R3(config)# interface Serial0/0/0
R3(config-if)# mpls label protocol ldp
R3(config-if)# mpls ip
R3(config-if)# interface Serial0/0/1
```

```
R3(config-if)# mpls label protocol ldp
R3(config-if)# mpls ip

R4(config)# mpls label protocol tdp
R4(config)# interface GigabitEthernet0/0
R4(config-if)# mpls label protocol ldp
R4(config-if)# mpls ip
R4(config-if)# interface Serial0/0/1
R4(config-if)# mpls label protocol ldp
R4(config-if)# mpls ip

R5(config)# mpls label protocol tdp
R5(config)# interface GigabitEthernet0/0
R5(config-if)# mpls label protocol ldp
R5(config-if)# mpls ip
R5(config-if)# interface Serial0/0/1
R5(config-if)# mpls label protocol ldp
R5(config-if)# mpls ip

R6(config)# mpls label protocol tdp
R6(config)# interface GigabitEthernet0/0
R6(config-if)# mpls label protocol ldp
R6(config-if)# mpls ip
```

Example 3-6 shows verification of the configuration with the LDP peering between each router. Notice that the loop-back addresses are used for LDP peer identification.

EXAMPLE 3-6 MPLS Configuration Verification

```
R1# show mpls ldp neighbor
Peer LDP Ident: 120.100.2.1:0; Local LDP Ident 120.100.1.1:0
TCP connection: 120.100.2.1.40418 - 120.100.1.1.646
State: Oper; Msgs sent/rcvd: 69/71; Downstream
Up time: 00:47:20
LDP discovery sources:
  Serial0/0/0, Src IP addr: 120.100.123.2
Addresses bound to peer LDP Ident:
  120.100.123.2 120.100.25.2 120.100.2.1
Peer LDP Ident: 120.100.3.1:0; Local LDP Ident 120.100.1.1:0
```

```
TCP connection: 120.100.3.1.51369 - 120.100.1.1.646
State: Oper; Msgs sent/rcvd: 68/68; Downstream
Up time: 00:47:18
LDP discovery sources:
  Serial0/0/0, Src IP addr: 120.100.123.3
Addresses bound to peer LDP Ident:
  120.100.123.3  120.100.3.1  120.100.34.3
```

R2# show mpls ldp neighbor

```
Peer LDP Ident: 120.100.3.1:0; Local LDP Ident 120.100.2.1:0
TCP connection: 120.100.3.1.16991 - 120.100.2.1.646
State: Oper; Msgs sent/rcvd: 71/68; Downstream
Up time: 00:46:33
LDP discovery sources:
  Serial0/0, Src IP addr: 120.100.123.3
  Serial0/1, Src IP addr: 120.100.34.3
Addresses bound to peer LDP Ident:
  120.100.123.3  120.100.3.1  120.100.34.3
Peer LDP Ident: 120.100.5.1:0; Local LDP Ident 120.100.2.1:0
TCP connection: 120.100.5.1.13826 - 120.100.2.1.646
State: Oper; Msgs sent/rcvd: 73/76; Downstream
Up time: 00:46:24
LDP discovery sources:
  Serial0/1, Src IP addr: 120.100.25.5
Addresses bound to peer LDP Ident:
  120.100.25.5  120.100.5.1  5.5.5.5  120.100.45.5
  100.100.100.5
Peer LDP Ident: 120.100.1.1:0; Local LDP Ident 120.100.2.1:0
TCP connection: 120.100.1.1.646 - 120.100.2.1.40418
State: Oper; Msgs sent/rcvd: 69/68; Downstream
Up time: 00:46:07
LDP discovery sources:
  Serial0/0, Src IP addr: 120.100.123.1
Addresses bound to peer LDP Ident:
  120.100.123.1  120.100.1.1
Peer LDP Ident: 120.100.4.1:0; Local LDP Ident 120.100.2.1:0
TCP connection: 120.100.4.1.47401 - 120.100.2.1.646
State: Oper; Msgs sent/rcvd: 54/57; Downstream
Up time: 00:32:28
LDP discovery sources:
```

```
Serial0/1, Src IP addr: 120.100.34.4
Addresses bound to peer LDP Ident:
120.100.4.1      4.4.4.4      120.100.45.4   100.100.100.4
120.100.34.4
```

R3# **show mpls ldp neighbor**

```
Peer LDP Ident: 120.100.2.1:0; Local LDP Ident 120.100.3.1:0
TCP connection: 120.100.2.1.646 - 120.100.3.1.16991
State: Oper; Msgs sent/rcvd: 69/72; Downstream
Up time: 00:47:11
LDP discovery sources:
  Serial0/0/0, Src IP addr: 120.100.123.2
  Serial0/0/1, Src IP addr: 120.100.25.2
Addresses bound to peer LDP Ident:
  120.100.123.2  120.100.25.2  120.100.2.1
Peer LDP Ident: 120.100.1.1:0; Local LDP Ident 120.100.3.1:0
TCP connection: 120.100.1.1.646 - 120.100.3.1.51369
State: Oper; Msgs sent/rcvd: 67/67; Downstream
Up time: 00:46:43
LDP discovery sources:
  Serial0/0/0, Src IP addr: 120.100.123.1
Addresses bound to peer LDP Ident:
  120.100.123.1  120.100.1.1
Peer LDP Ident: 120.100.5.1:0; Local LDP Ident 120.100.3.1:0
TCP connection: 120.100.5.1.53107 - 120.100.3.1.646
State: Oper; Msgs sent/rcvd: 67/74; Downstream
Up time: 00:45:22
LDP discovery sources:
  Serial0/0/1, Src IP addr: 120.100.25.5
Addresses bound to peer LDP Ident:
  120.100.25.5  120.100.5.1  5.5.5.5  120.100.45.5
  100.100.100.5
Peer LDP Ident: 120.100.4.1:0; Local LDP Ident 120.100.3.1:0
TCP connection: 120.100.4.1.15940 - 120.100.3.1.646
State: Oper; Msgs sent/rcvd: 52/56; Downstream
Up time: 00:33:06
LDP discovery sources:
  Serial0/0/1, Src IP addr: 120.100.34.4
Addresses bound to peer LDP Ident:
  120.100.4.1  4.4.4.4  120.100.45.4  100.100.100.4
```

120.100.34.4

R4# **show mpls ldp neighbor**

```
Peer LDP Ident: 120.100.6.1:0; Local LDP Ident 120.100.4.1:0
  TCP connection: 120.100.6.1.55234 - 120.100.4.1.646
  State: Oper; Msgs sent/rcvd: 74/76; Downstream
  Up time: 00:43:52
  LDP discovery sources:
    GigabitEthernet0/0, Src IP addr: 120.100.45.6
  Addresses bound to peer LDP Ident:
    120.100.6.1    6.6.6.6    100.100.100.6    120.100.45.6
Peer LDP Ident: 120.100.5.1:0; Local LDP Ident 120.100.4.1:0
  TCP connection: 120.100.5.1.57689 - 120.100.4.1.646
  State: Oper; Msgs sent/rcvd: 72/74; Downstream
  Up time: 00:43:48
  LDP discovery sources:
    GigabitEthernet0/0, Src IP addr: 120.100.45.5
    Serial0/0/1, Src IP addr: 120.100.25.5
  Addresses bound to peer LDP Ident:
    120.100.25.5    120.100.5.1    5.5.5.5    120.100.45.5
    100.100.100.5
Peer LDP Ident: 120.100.2.1:0; Local LDP Ident 120.100.4.1:0
  TCP connection: 120.100.2.1.646 - 120.100.4.1.47401
  State: Oper; Msgs sent/rcvd: 55/52; Downstream
  Up time: 00:30:52
  LDP discovery sources:
    Serial0/0/1, Src IP addr: 120.100.25.2
  Addresses bound to peer LDP Ident:
    120.100.123.2    120.100.25.2    120.100.2.1
Peer LDP Ident: 120.100.3.1:0; Local LDP Ident 120.100.4.1:0
  TCP connection: 120.100.3.1.646 - 120.100.4.1.15940
  State: Oper; Msgs sent/rcvd: 54/50; Downstream
  Up time: 00:30:52
  LDP discovery sources:
    Serial0/0/1, Src IP addr: 120.100.34.3
  Addresses bound to peer LDP Ident:
    120.100.123.3    120.100.3.1    120.100.34.3
```

R5# **show mpls ldp neighbor**

```
Peer LDP Ident: 120.100.2.1:0; Local LDP Ident 120.100.5.1:0
```

```
TCP connection: 120.100.2.1.646 - 120.100.5.1.13826
State: Oper; Msgs sent/rcvd: 80/77; Downstream
Up time: 00:49:55
LDP discovery sources:
  Serial0/0/1, Src IP addr: 120.100.25.2
Addresses bound to peer LDP Ident:
  120.100.123.2  120.100.25.2  120.100.2.1
Peer LDP Ident: 120.100.6.1:0; Local LDP Ident 120.100.5.1:0
TCP connection: 120.100.6.1.18472 - 120.100.5.1.646
State: Oper; Msgs sent/rcvd: 81/81; Downstream
Up time: 00:48:58
LDP discovery sources:
  GigabitEthernet0/0, Src IP addr: 120.100.45.6
Addresses bound to peer LDP Ident:
  120.100.6.1  6.6.6.6  100.100.100.6  120.100.45.6
Peer LDP Ident: 120.100.4.1:0; Local LDP Ident 120.100.5.1:0
TCP connection: 120.100.4.1.646 - 120.100.5.1.57689
State: Oper; Msgs sent/rcvd: 80/78; Downstream
Up time: 00:48:54
LDP discovery sources:
  GigabitEthernet0/0, Src IP addr: 120.100.45.4
  Serial0/0/1, Src IP addr: 120.100.34.4
Addresses bound to peer LDP Ident:
  120.100.4.1  4.4.4.4  120.100.45.4  100.100.100.4
  120.100.34.4
Peer LDP Ident: 120.100.3.1:0; Local LDP Ident 120.100.5.1:0
TCP connection: 120.100.3.1.646 - 120.100.5.1.53107
State: Oper; Msgs sent/rcvd: 77/70; Downstream
Up time: 00:48:17
LDP discovery sources:
  Serial0/0/1, Src IP addr: 120.100.34.3
Addresses bound to peer LDP Ident:
  120.100.123.3  120.100.3.1  120.100.34.3
```

R6# show mpls ldp neighbor

```
Peer LDP Ident: 120.100.5.1:0; Local LDP Ident 120.100.6.1:0
TCP connection: 120.100.5.1.646 - 120.100.6.1.18472
State: Oper; Msgs sent/rcvd: 82/82; Downstream
Up time: 00:49:31
LDP discovery sources:
```

```

GigabitEthernet0/0, Src IP addr: 120.100.45.5
Addresses bound to peer LDP Ident:
 120.100.25.5    120.100.5.1    5.5.5.5        120.100.45.5
 100.100.100.5
Peer LDP Ident: 120.100.4.1:0; Local LDP Ident 120.100.6.1:0
TCP connection: 120.100.4.1.646 - 120.100.6.1.55234
State: Oper; Msgs sent/rcvd: 82/80; Downstream
Up time: 00:49:31
LDP discovery sources:
  GigabitEthernet0/0, Src IP addr: 120.100.45.4
Addresses bound to peer LDP Ident:
 120.100.4.1    4.4.4.4        120.100.45.4    100.100.100.4
 120.100.34.4

```

- You will be configuring two VPNs over your MPLS networks per Figure 3-9 between PE routers of BLUE and RED. At this point, assign the following interfaces on each PE router into separate routing instances within the routers:

PE R1 interface Gi0/0 VLAN10 connection into VPN BLUE

PE R1 interface Gi0/0 VLAN50 connection into VPN RED PE

R6 interface Gi0/1 VLAN20 connection into VPN BLUE PE

R6 interface Gi0/1 VLAN100 connection into VPN RED

Configure VPN BLUE to use an RD of 100 and VPN RED to use an RD of 200 for both importing and exporting routes into your BGP network, which will be configured later with an AS of AS65001. (4 points)

You are required to create virtual routing forwarding (VRF) instances on the PE routers and assign the subinterfaces on each PE router into these. This will ultimately provide end-to-end virtual private networking (VPN) connectivity over the MPLS network for your CE devices to communicate. You are directed to use a route descriptor (RD) of 100 for the BLUE VRF and 200 for the RED VRF and must combine this with the BGP autonomous system (AS) number of 65001 to import and export route target extended communities for the specified VRFs. The actual BGP configuration will be configured later in the lab. If you have configured this correctly, as shown in Example 3-7, you have scored 4 points.

EXAMPLE 3-7 VRF Configuration

```

R1(config)# ip vrf BLUE
R1(config-vrf)# rd 65001:100
R1(config-vrf)# route-target export 65001:100

```

```
R1(config-vrf)# route-target import 65001:100
R1(config-vrf)#!R1(config-
vrf)# ip vrf RED R1(config-
vrf)# rd 65001:200
R1(config-vrf)# route-target export 65001:200
R1(config-vrf)# route-target import 65001:200
R1(config-vrf)# exit
R1(config)# interface GigabitEthernet0/0.10
R1(config-subif)# ip vrf forwarding BLUE
R1(config-subif)# interface GigabitEthernet0/0.50
R1(config-subif)# ip vrf forwarding RED
```

```
R6(config)# ip vrf BLUE
```

```
R6(config-vrf)# rd 65001:100
```

```
R6(config-vrf)# route-target export 65001:100
```

```
R6(config-vrf)# route-target import 65001:100
```

```
R6(config-vrf)# ip vrf RED
```

```
R6(config-vrf)# rd 65001:200
```

```
R6(config-vrf)# route-target export 65001:200
```

```
R6(config-vrf)# route-target import 65001:200
```

```
R6(config-vrf)# exit
```

```
R6(config)# interface GigabitEthernet0/1.20
```

```
R6(config-subif)# ip vrf forwarding BLUE
```

```
R6(config)# interface GigabitEthernet0/1.100 ip vrf forwarding RED
```

- Create a network between PE Router R1 and CE device Sw1 using a VLAN10 interface on Sw1 that can be trunked toward R1. This network will reside in the BLUE VPN. Use a subnet of 10.10.10.0/30 with .1/30 assigned to the PE and .2/30 assigned to the CE. (2 points)

This is a simple configuration task to assign IP connectivity between the PE and CE devices for future routing between the devices and remote VPN connectivity via R6. The new VLAN10 must be created on Sw1, and this VLAN should have already been permitted to flow through to R1 as an allowed VLAN. The subinterface of Gigabit0/0.10 on R1 has been assigned to the BLUE VRF during the previous question, so connectivity between Sw1 and R1 should now be possible (when IP addresses are assigned). When testing, remember that R1 must use the appropriate VRF to confirm

connectivity because a normal ping would be sourced from the global routing table and will fail. If you have configured this correctly, as shown in Example 3-8, you have scored 2 points.

EXAMPLE 3-8 BLUE VRF IP Addressing and Local Connectivity Testing

```
R1(config)# interface GigabitEthernet0/0.10
R1(config-subif)# ip add 10.10.10.1 255.255.255.252

Switch1(config)# vlan 10
Switch1(config-vlan)# exit
Switch1(config)# interface vlan 10
Switch1(config-if)# no shutdown
Switch1(config-if)# ip add 10.10.10.2 255.255.255.252

R1# ping vrf BLUE 10.10.10.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.10.2, timeout is 2 seconds:
..!!!
Success rate is 60 percent (3/5), round-trip min/avg/max = 1/1/1 ms
```

- Create a network between PE router R6 and CE device Sw2 using a VLAN20 interface on Sw2 that can be trunked toward R6. This network will reside in the BLUE VPN. Use a subnet of 10.10.20.0/30 with .1/30 assigned to the PE and .2/30 assigned to the CE. (2 points)

This is a simple configuration task as per the previous question to assign connectivity between the PE and CE devices for future routing between the devices and remote VPN connectivity via R1. The new VLAN20 must be created on Sw2, and this VLAN already should have been permitted to flow through to R6 as an allowed VLAN. The subinterface of Gigabit0/1.20 on R6 has been assigned to the BLUE VRF during a previous question, so connectivity between Sw2 and R6 should now be possible. When testing, remember that R6 must use the appropriate VRF to confirm connectivity. If you have configured this correctly, as shown in Example 3-9, you have scored 2 points.

EXAMPLE 3-9 BLUE VRF IP Addressing and Local Connectivity Testing

```
R6(config)# interface GigabitEthernet0/1.20
R6(config-subif)# ip add 10.10.20.1 255.255.255.252

Switch2(config)# vlan 20
Switch2(config-vlan)# exit
```

```
Switch2(config)# interface vlan 20
Switch2(config-if)# no shutdown
Switch2(config-if)# ip add 10.10.20.2 255.255.255.252

R6# ping vrf BLUE 10.10.20.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.20.2, timeout is 2 seconds:
..!!!
Success rate is 60 percent (3/5), round-trip min/avg/max = 1/1/1 ms
```

- Create a network between PE Router R1 and CE device Sw3 using a VLAN50 interface on Sw3 that can be trunked toward R1; this network will reside in the RED VPN. Use a subnet of 130.50.50.0/30 with .1/30 assigned to the PE and .2/30 assigned to the CE. (2 point)

Here's another simple configuration to assign connectivity between the PE and CE devices for future routing between the devices and remote VPN connectivity via R6. The new VLAN50 must be created on Sw3, and this VLAN should have already been permitted to flow through Sw1 to R1 as an allowed VLAN. The subinterface of Gigabit0/0.50 on R1 has been assigned to the RED VRF during a previous question, so connectivity between Sw3 and R1 should now be possible. When testing, remember that R1 must use the appropriate VRF to confirm connectivity. If you have configured this correctly, as shown in Example 3-10, you have scored 2 points.

EXAMPLE 3-10 RED VRF IP Addressing and Local Connectivity Testing

```
R1(config)# interface GigabitEthernet0/0.50
R1(config-subif)# ip add 130.50.50.1 255.255.255.252

Switch3(config)# vlan 50
Switch3(config-vlan)# exit
Switch3(config)# interface vlan 50
Switch3(config-if)# no shutdown
Switch3(config-if)# ip add 130.50.50.2 255.255.255.252

R1# ping vrf RED 130.50.50.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 130.50.50.2, timeout is 2 seconds:
```

```
..!!!
```

```
Success rate is 60 percent (3/5), round-trip min/avg/max = 1/1/1 ms
```

- Create a network between PE Router R6 and CE device Sw4 using a VLAN100 interface on Sw4 that can be trunked toward R6; this network will reside in the RED VPN. Use a subnet of 130.100.100.0/30 with .1/30 assigned to the PE and .2/30 assigned to the CE. (2 points)

This is the final configuration task to assign connectivity between the PE and CE devices for future routing between the devices and remote VPN connectivity via R1. The new VLAN100 must be created on Sw4, and this VLAN should have already been permitted to flow through Sw3 to R6 as an allowed VLAN. The subinterface of Gigabit0/1.100 on R6 has been assigned to the RED VRF during a previous question, so connectivity between Sw4 and R6 should now be possible. When testing, remember that R6 must use the appropriate VRF to confirm connectivity. If you have configured this correctly, as shown in Example 3-11, you have scored 2 points.

EXAMPLE 3-11 RED VRF IP Addressing and Local Connectivity Testing

```
R6(config)# interface GigabitEthernet0/1.100
R6(config-subif)# ip add 130.100.100.1 255.255.255.252
```

```
Switch4(config)# vlan 100
Switch4(config-vlan)# exit
Switch4(config)# interface vlan 100
Switch4(config-if)# no shutdown
Switch4(config-if)# ip add 130.100.100.2 255.255.255.252
```

```
R6# ping vrf RED 130.100.100.2
```

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 130.100.100.2, timeout is 2 seconds:
..!!!
Success rate is 60 percent (3/5), round-trip min/avg/max = 1/1/1 ms
```

Section 3: BGP (5 Points)

- Configure MP-BGP between your PE routers, per Figure 3-10, to enable your network to transport the VPNv4 addresses of your configured VPNs (BLUE and RED). Use loopback interfaces for peering between your PE routers. You will configure the actual VPN routing in later questions. (4 points)

MPLS requires the use of Multiprotocol BGP (MP-BGP) between the PE routers to exchange VPNv4 addresses in addition to IPv4 addresses. The VPNs will be mapped into the configuration later, so this question is a straightforward peering and VPNv4 setup task. The configuration requires you to peer from your loopback interfaces, which are advertised via your P routers within OSPF and that extended communities are used between PE routers to advertise your VPNv4 addresses successfully. You should be aware that Route Targets (RT) are implemented by the use of the BGP extended community (64 bits) and as such the **send-community both** value must be configured within MP-BGP. The **next-hop-self** command is optional and strictly required only when you have an eBGP configuration to preserve the next-hop information to peers; you won't lose any points if you added this or left it out. The actual VPN portion of MP-BGP will be configured later within the IPv4 address family for VRF-specific advertisements. This is a simple MP-BGP network with only two PE routers; additional PE routers would require a full mesh of iBGP peering or configuration of route-reflectors to aid scalability.

If you have configured this correctly, as shown in Example 3-12, you have scored 4 points.

EXAMPLE 3-12 MP-BGP Configuration

```
R1(config)# router bgp 65001
R1(config-router)# no synchronization
R1(config-router)# no auto-summary
R1(config-router)# neighbor 120.100.6.1 remote-as 65001
R1(config-router)# neighbor 120.100.6.1 update-source Loopback0
R1(config-router)# address-family vpnv4
R1(config-router-af)# neighbor 120.100.6.1 activate
R1(config-router-af)# neighbor 120.100.6.1 next-hop-self
R1(config-router-af)# neighbor 120.100.6.1 send-community both

R6(config)# router bgp 65001
R6(config-router)# no sync
R6(config-router)# no auto-summary
R6(config-router)# neighbor 120.100.1.1 remote-as 65001
R6(config-router)# neighbor 120.100.1.1 update-source Loopback0

R6(config-router)# address-family vpnv4
R6(config-router-af)# neighbor 120.100.1.1 activate
R6(config-router-af)# neighbor 120.100.1.1 next-hop-self
R6(config-router-af)# neighbor 120.100.1.1 send-community both
```

Section 4: EIGRP and MP-BGP (9 Points)

- Configure EIGRP per Figure 3-11 between your PE Router R6 and CE Switch Sw2. Use an EIGRP process number of 1 on R6 and a process number of 10 on Sw2. Use VLAN20 for EIGRP connectivity between R6 and Sw2. Advertise all preconfigured Loopback networks on Sw2 to R6 for the BLUE VPN. (3 points)

Until now the questions have merely dealt with setting up the infrastructure for MPLS connectivity. Now you are requested to advertise routes from your CE Switch Sw2 to PE Router R6, which will ultimately be advertised throughout the BLUE VPN to the remote PE Router R1 and CE Switch Sw1. The questions become harder from this point. You'll realize that to peer successfully with EIGRP you would need to be operating within the same autonomous system (AS) number, yet the question enforces you to run differing AS numbers. PE routers would normally connect to multiple customers, so it is unreasonable to expect that each EIGRP domain should run the same AS number. As such, there is a fix, which is a manual AS mapping under the VPN-specific configuration (**address-family ipv4 vrf BLUE**) where the AS number is stipulated. It is also within this section that the networks are enabled for EIGRP to operate over. Example 3-13 details the EIGRP configuration and resulting neighbor relationship and route propagation between R6 and Sw2. If you have configured this correctly, as shown in Example 3-13, you have scored 3 points.

EXAMPLE 3-13 R6 and Switch2 EIGRP Configuration and Verification

NOTE

The IP addressing for VLAN20 on Sw2 and associated subinterfaces on R6 has previously been configured. The BLUE VRF has also been associated to the R6 subinterface previously.

```
R6(config)# router eigrp 1
R6(config-router)# address-family ipv4 vrf BLUE
R6(config-router-af)# autonomous-system 10
R6(config-router-af)# no auto-summary
R6(config-router-af)# network 10.10.20.0 0.0.0.3

Switch2(config)# ip routing
Switch2(config)# router eigrp 10
Switch2(config-router)# no auto-summary
Switch2(config-router)# network 10.10.20.0 0.0.0.3
Switch2(config-router)# network 10.2.2.0 0.0.0.255
Switch2(config-router)# network 10.2.3.0 0.0.0.255
Switch2(config-router)# network 10.2.4.0 0.0.0.255

R6# show ip eigrp vrf BLUE neighbors
IP-EIGRP neighbors for process 10
H   Address                               Interface      Hold Uptime    SRTT    RTO    Q    Seq
```



```

                                (sec)      (ms)      Cnt Num
0  10.10.20.2                    Gi0/1.20      11 00:04:18    1  200  0  1
R6#
R6# show ip route vrf BLUE eigrp
      10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
D      10.2.2.0/24
          [90/156160] via 10.10.20.2, 00:04:36, GigabitEthernet0/1.20
D      10.2.3.0/24
          [90/156160] via 10.10.20.2, 00:04:36, GigabitEthernet0/1.20
D      10.2.4.0/24
          [90/156160] via 10.10.20.2, 00:04:36, GigabitEthernet0/1.20

```

- Configure EIGRP per Figure 3-11 between your PE Router R1 and CE Switch Sw1. Use an EIGRP process number of 1 on R1 and a process number of 10 on Sw1. Use VLAN10 for EIGRP connectivity between R1 and Sw1. Advertise all preconfigured Loopback networks on Sw1 to R1 for the BLUE VPN. (3 points)

NOTE

The IP addressing for VLAN10 on Sw1 and associated subinterfaces on R1 has previously been configured. The BLUE VRF has also been associated to the R1 subinterface previously.

Per the previous question, you are requested to advertise routes from your CE Switch Sw1 to PE Router R1, which will ultimately be advertised throughout the BLUE VPN to the remote PE Router R6 and CE Switch Sw2. Once again you are required to manually configure the EIGRP AS number within the address-family vrf section of the PE. Example 3-14 details the EIGRP configuration and resulting neighbor relationship and route propagation between R1 and Sw1. If you have configured this correctly, as shown in Example 3-14, you have scored 3 points.

EXAMPLE 3-14 R1 and Switch1 EIGRP Configuration and Verification

```

R1(config)# router eigrp 1
R1(config-router)# address-family ipv4 vrf BLUE
R1(config-router-af)# autonomous-system 10
R1(config-router-af)# no auto-summary
R1(config-router-af)# network 10.10.10.0 0.0.0.3
R1(config-vrf)# int gi0/0.10
R1(config-subif)# ip vrf forwarding BLUE
R1(config-subif)# ip add 10.10.10.1 255.255.255.252

Switch1(config)# ip routing
Switch1(config)# router eigrp 10
Switch1(config-router)# no auto-summary
Switch1(config-router)# network 10.10.10.0 0.0.0.3
Switch1(config-router)# network 10.1.1.0 0.0.0.255
Switch1(config-router)# network 10.1.2.0 0.0.0.255

```

```

Switch1(config-router)# network 10.1.3.0 0.0.0.255

R1# show ip eigrp vrf BLUE neighbors
IP-EIGRP neighbors for process 10
H   Address                Interface          Hold Uptime    SRTT   RTO   Q   Seq
                               (sec)           (ms)          Cnt  Num
0   10.10.10.2              Gi0/0.10          13 00:00:24    1    200  0   1
R1#

R1# show ip eigrp vrf BLUE neighbors
IP-EIGRP neighbors for process 10
H   Address                Interface          Hold Uptime    SRTT   RTO   Q   Seq
                               (sec)           (ms)          Cnt  Num
0   10.10.10.2              Gi0/0.10          13 00:00:24    1    200  0   1
R1# show ip route vrf BLUE eigrp
10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
D    10.1.3.0/24
      [90/153856] via 10.10.10.2, 00:01:18, GigabitEthernet0/0.10
D    10.1.2.0/24
      [90/153856] via 10.10.10.2, 00:01:18, GigabitEthernet0/0.10
D    10.1.1.0/24
      [90/153856] via 10.10.10.2, 00:01:18, GigabitEthernet0/0.10

```

- Configure your PE Routers R1 and R6 to transport EIGRP routes from your CE devices between the BLUE VPN using MP-BGP. EIGRP networks residing on Sw1 should be seen as internal EIGRP routes on Sw2 and vice versa. Ensure all EIGRP routes have a MED of 50 assigned to them within MP-BGP. Use a default-metric of 10000 100 255 1 1500 for BGP routes when redistributed into EIGRP. (3 points)

The full end-to-end VPN routing is achieved at this point by redistributing EIGRP into the appropriate address-family for the VRF. The question dictates the metrics you should use. In reality, the metrics are not required because the extended community values of MP-BGP previously configured will effectively transport the internal metrics of EIGRP and ensure the routes are shown as internal EIGRP routes at the remote location, even though they have been redistributed via another routing protocol. The question is just looking for accuracy and giving you the opportunity to view routes with the metrics and later without if you choose to. Example 3-15 details the configuration required on the PE routers and resulting routes on the CE devices Sw1 and Sw2. If you have configured this correctly, as shown in Example 3-15, you have scored 3 points.

EXAMPLE 3-15 PE and CE MP-BGP Redistribution Configuration and Verification

```

R1(config)# router eigrp 1
R1(config-router)# address-family ipv4 vrf BLUE
R1(config-router-af)# redistribute bgp 65001 metric 10000 100 255 1 1500
R1(config-router-af)# router bgp 65001
R1(config-router)# address-family ipv4 vrf BLUE
R1(config-router-af)# redistribute eigrp 10 metric 50

R6(config)# router eigrp 1
R6(config-router)# address-family ipv4 vrf BLUE
R6(config-router-af)# redistribute bgp 65001 metric 10000 100 255 1 1500
R6(config-router-af)# router bgp 65001
R6(config-router)# address-family ipv4 vrf BLUE
R6(config-router-af)# redistribute eigrp 10 metric 50

SW1# show ip route eigrp
D 10.2.2.0/24 [90/156416] via 10.10.10.1, 00:32:05, Vlan10
D 10.2.3.0/24 [90/156416] via 10.10.10.1, 00:32:05, Vlan10
D 10.2.4.0/24 [90/156416] via 10.10.10.1, 00:32:05, Vlan10
D 10.10.20.0/30 [90/28416] via 10.10.10.1, 00:32:05, Vlan10

SW2# show ip route eigrp
D 10.1.3.0/24 [90/154112] via 10.10.20.1, 00:33:07, Vlan20
D 10.1.2.0/24 [90/154112] via 10.10.20.1, 00:33:07, Vlan20
D 10.1.1.0/24 [90/154112] via 10.10.20.1, 00:33:07, Vlan20
D 10.10.10.0/30 [90/26112] via 10.10.20.1, 00:33:07, Vlan20

```

Example 3-16 details the BGP routes received on the PE routers with the assigned MED value of 50; it also details the MPLS forwarding table for the BLUE VRF. Notice the iBGP routes on the PE routers from the remote PE router with the MED of 50; these are the routes that are propagated to EIGRP CE devices. If you have configured this correctly, as shown in Example 3-16, you have scored 3 points.

EXAMPLE 3-16 PE MP-BGP and MPLS Verification

```

R6# show ip bgp vpnv4 vrf BLUE
BGP table version is 17, local router ID is 120.100.6.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale

```

Origin codes: i - IGP, e - EGP, ? - incomplete

```

      Network          Next Hop          Metric LocPrf Weight Path
Route Distinguisher: 65001:100 (default for vrf BLUE)
*>i10.1.1.0/24        120.100.1.1          50      100      0 ?
*>i10.1.2.0/24        120.100.1.1          50      100      0 ?
*>i10.1.3.0/24        120.100.1.1          50      100      0 ?
*> 10.2.2.0/24        10.10.20.2           50              32768 ?
*> 10.2.3.0/24        10.10.20.2           50              32768 ?
*> 10.2.4.0/24        10.10.20.2           50              32768 ?
*>i10.10.10.0/30     120.100.1.1           0      100      0 ?
*> 10.10.20.0/30     0.0.0.0               0              32768 ?

```

R1# **show ip bgp vpnv4 vrf BLUE**

BGP table version is 17, local router ID is 120.100.1.1

Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
r RIB-failure, S Stale

Origin codes: i - IGP, e - EGP, ? - incomplete

```

      Network          Next Hop          Metric LocPrf Weight Path
Route Distinguisher: 65001:100 (default for vrf BLUE)
*> 10.1.1.0/24        10.10.10.2           50              32768 ?
*> 10.1.2.0/24        10.10.10.2           50              32768 ?
*> 10.1.3.0/24        10.10.10.2           50              32768 ?
*>i10.2.2.0/24        120.100.6.1          50      100      0 ?
*>i10.2.3.0/24        120.100.6.1          50      100      0 ?
*>i10.2.4.0/24        120.100.6.1          50      100      0 ?
*> 10.10.10.0/30     0.0.0.0               0              32768 ?
*>i10.10.20.0/30     120.100.6.1           0      100      0 ?

```

R1# **show mpls forwarding-table vrf BLUE**

Local tag	Outgoing tag or VC	Prefix or Tunnel Id	Bytes switched	tag	Outgoing interface	Next Hop
26	Untagged	10.1.3.0/24 [V]	0		Gi0/0.10	10.10.10.2
27	Untagged	10.1.2.0/24 [V]	0		Gi0/0.10	10.10.10.2
28	Aggregate	10.10.10.0/30 [V]	0			
29	Untagged	10.1.1.0/24 [V]	0		Gi0/0.10	10.10.10.2

R6# **show mpls forwarding-table vrf BLUE**

Local tag	Outgoing tag or VC	Prefix or Tunnel Id	Bytes switched	tag	Outgoing interface	Next Hop
26	Untagged	10.2.2.0/24 [V]	0		Gi0/1.20	10.10.20.2
27	Untagged	10.2.3.0/24 [V]	0		Gi0/1.20	10.10.20.2
28	Untagged	10.2.4.0/24 [V]	0		Gi0/1.20	10.10.20.2
29	Aggregate	10.10.20.0/30 [V]	0			

Section 5: OSPF and MP-BGP (9 Points)

- Configure OSPF per Figure 3-12 for your VRF RED with a process number of 3 on PE Router R1 and Sw3 using VLAN50 for connectivity. Use a process ID of 2 on PE Router R6 and CE device Sw4 using VLAN100 for connectivity. You should permit only internal OSPF routes to be advertised across your VPN and ensure the redistribution of BGP routes into OSPF are assigned as Type 1 external routes with no manually adjusted cost associated to them. It is acceptable for these routes to come through as /32 routes because of default OSPF behavior of Loopback interfaces. (3 points)

NOTE

The IP addressing for VLAN50 on Sw3 and associated subinterface on R1 and VLAN100 on Sw4 and associated subinterface on R6 has previously been configured. The RED VRF has also been associated to the R1 and R6 subinterfaces previously.

You are requested to configure OSPF over your MPLS network between CE devices Sw3 and Sw4 via your PE Routers R1 and R6. Figure 3-12 indicates that all loopback interfaces are to be included in OSPF on both CE devices. You should be aware that OSPF will advertise these as host routes, but the question states that this is acceptable behavior. Similarly to the EIGRP question, you are requested to manipulate the redistribution of the IGP into BGP, but in reality the routes would appear to have not been redistributed through another routing protocol by default. This direction is actually a red herring for the next question when the routes at the CE devices appear as external routes when they should in fact be internal routes. You are requested to permit only internal OSPF routes to be redistributed into BGP, which is a simple **match internal** parameter on the redistribution configuration. You should, of course, remember that the MPLS network is seen as an OSPF super backbone, and as such you had no configuration for Area 0 to enable Area 1 to communicate with Area 2 over MPLS. Example 3-17 details the required configuration and verification. If you have configured this correctly, as shown in Example 3-17, you have scored 3 points.

EXAMPLE 3-17 VRF RED OSPF Configuration and Verification

```
SW3(config)# ip routing
SW3(config)# router ospf 3
SW3(config-router)# network 130.50.50.0 0.0.0.3 area 0
SW3(config-router)# network 10.33.33.0 0.0.0.255 area 1
SW3(config-router)# network 10.33.34.0 0.0.0.255 area 1
SW3(config-router)# network 10.33.35.0 0.0.0.255 area 1
```

```
SW4(config)# ip routing
SW4(config)# router ospf 2
SW4(config-router)# network 130.100.100.0 0.0.0.3 area 0
SW4(config-router)# network 10.44.44.0 0.0.0.255 area 2
SW4(config-router)# network 10.44.45.0 0.0.0.255 area 2
SW4(config-router)# network 10.44.46.0 0.0.0.255 area 2
```

```
R1(config)# router ospf 3 vrf RED
R1(config-router)# network 130.50.50.0 0.0.0.3 area 0
R1(config-router)# redistribute bgp 65001 subnets metric-type 1
R1(config-router)# router bgp 65001
R1(config-router)# address-family ipv4 vrf RED
R1(config-router-af)# redistribute ospf 3 match internal
```

```
R6(config)# router ospf 2 vrf RED
R6(config-router)# net 130.100.100.0 0.0.0.3 area 0
R6(config-router)# redistribute bgp 65001 subnets metric-type 1
R6(config-router)# router bgp 65001
R6(config-router)# address-family ipv4 vrf RED
R6(config-router-af)# redistribute ospf 2 match internal
```

```
R1# show ip route vrf RED ospf
```

```
Routing Table: RED
```

```
10.0.0.0/32 is subnetted, 6 subnets
O IA 10.33.34.1 [110/2] via 130.50.50.2, 00:04:48, GigabitEthernet0/0.50
O IA 10.33.35.1 [110/2] via 130.50.50.2, 00:04:48, GigabitEthernet0/0.50
O IA 10.33.33.1 [110/2] via 130.50.50.2, 00:04:48, GigabitEthernet0/0.50
```

```
R6# show ip route vrf RED ospf
```

```
Routing Table: RED
```

```
10.0.0.0/32 is subnetted, 6 subnets
O IA 10.44.46.1 [110/2] via 130.100.100.2, 00:02:32, GigabitEthernet0/1.100
O IA 10.44.45.1 [110/2] via 130.100.100.2, 00:02:32, GigabitEthernet0/1.100
O IA 10.44.44.1 [110/2] via 130.100.100.2, 00:02:32, GigabitEthernet0/1.100
```

```

SW3# show ip route ospf
    130.100.0.0/30 is subnetted, 1 subnets
O E1   130.100.100.0 [110/2] via 130.50.50.1, 00:06:08, Vlan50
    10.0.0.0/8 is variably subnetted, 6 subnets, 2 masks
O E1   10.44.46.1/32 [110/3] via 130.50.50.1, 00:02:54, Vlan50
O E1   10.44.45.1/32 [110/3] via 130.50.50.1, 00:02:54, Vlan50
O E1   10.44.44.1/32 [110/3] via 130.50.50.1, 00:02:55, Vlan50

SW4# show ip route ospf
    130.50.0.0/30 is subnetted, 1 subnets
O E1   130.50.50.0 [110/2] via 130.100.100.1, 00:03:37, Vlan100
    10.0.0.0/8 is variably subnetted, 6 subnets, 2 masks
O E1   10.33.34.1/32 [110/3] via 130.100.100.1, 00:03:37, Vlan100
O E1   10.33.35.1/32 [110/3] via 130.100.100.1, 00:03:37, Vlan100
O E1   10.33.33.1/32 [110/3] via 130.100.100.1, 00:03:37, Vlan100

```

- You will notice that your OSPF IA (intra-area) routes between CE devices Sw3 and Sw4 appear as Type 1 External routes; configure your OSPF network appropriately to ensure the routes are displayed correctly as IA routes. You are not permitted to adjust the OSPF redistribution into BGP as directed in the previous question. Maintain the OSPF process IDs as previously directed, and you are permitted to configure only Router R1. (6 points)

This is a tricky question and one that will really eat into your time—the kind of question that if the answer doesn’t jump out at you and the points don’t look appealing enough, it’s one to park and come back to. You can leave questions like this confidently because you have your routes in place and following questions don’t build from this one. As stated previously, the redistribution into Type 1 is actually somewhat misleading. When you look at the routes in Example 3-18 for the PE routers, you will see that they are actually IA routes at this point, so it is only when these routes are advertised to the CE devices that the Type 1 External route change occurs.

EXAMPLE 3-18 VRF RED OSPF Routes

```

R1# show ip route vrf RED ospf

Routing Table: RED

    10.0.0.0/32 is subnetted, 6 subnets
O IA   10.33.34.1 [110/2] via 130.50.50.2, 00:04:48, GigabitEthernet0/0.50
O IA   10.33.35.1 [110/2] via 130.50.50.2, 00:04:48, GigabitEthernet0/0.50
O IA   10.33.33.1 [110/2] via 130.50.50.2, 00:04:48, GigabitEthernet0/0.50

```

```
R6# show ip route vrf RED ospf
```

```
Routing Table: RED
```

```
10.0.0.0/32 is subnetted, 6 subnets
O IA 10.44.46.1 [110/2] via 130.100.100.2, 00:02:32, GigabitEthernet0/1.100
O IA 10.44.45.1 [110/2] via 130.100.100.2, 00:02:32, GigabitEthernet0/1.100
O IA 10.44.44.1 [110/2] via 130.100.100.2, 00:02:32, GigabitEthernet0/1.100
```

```
SW3# show ip route ospf
```

```
130.100.0.0/30 is subnetted, 1 subnets
O E1 130.100.100.0 [110/2] via 130.50.50.1, 00:06:08, Vlan50
10.0.0.0/8 is variably subnetted, 6 subnets, 2 masks
O E1 10.44.46.1/32 [110/3] via 130.50.50.1, 00:02:54, Vlan50
O E1 10.44.45.1/32 [110/3] via 130.50.50.1, 00:02:54, Vlan50
O E1 10.44.44.1/32 [110/3] via 130.50.50.1, 00:02:55, Vlan50
```

```
SW4# show ip route ospf
```

```
130.50.0.0/30 is subnetted, 1 subnets
O E1 130.50.50.0 [110/2] via 130.100.100.1, 00:03:37, Vlan100
10.0.0.0/8 is variably subnetted, 6 subnets, 2 masks
O E1 10.33.34.1/32 [110/3] via 130.100.100.1, 00:03:37, Vlan100
O E1 10.33.35.1/32 [110/3] via 130.100.100.1, 00:03:37, Vlan100
O E1 10.33.33.1/32 [110/3] via 130.100.100.1, 00:03:37, Vlan100
```

The clue is actually in the question “Maintain the OSPF process IDs as previously directed.” Statements like this should make you think, “Okay, so if I did change the process ID, it would most likely work; why would that do it and how else can I achieve that?” OSPF has a domain ID; by default, this is the same as the process ID. If the process IDs are different on PE routers that form the VPN, the LSA is changed to a type 5 and the routes become external. You might not have known that, but it’s the kind of thing that you gain through research and rack time. Because you are not permitted to change the process ID, you are only left with the option of changing the domain ID. Example 3-19 details the domain ID information on your PE routers, the configuration required to change the domain ID on one of your PE’s Router R1, and the resulting IA routes received on your CE devices. If you have configured this correctly, as shown in Example 3-19, you have scored 6 points.

EXAMPLE 3-19 Domain ID Configuration and OSPF Route Verification

```
R1# show ip ospf 3 | include Domain
    Domain ID type 0x0005, value 0.0.0.3

R6# show ip ospf 2 | include Domain
    Domain ID type 0x0005, value 0.0.0.2

R1(config)# router ospf 3 vrf RED
R1(config-router)# domain-id 0.0.0.2

SW3# show ip route ospf
    130.100.0.0/30 is subnetted, 1 subnets
O IA   130.100.100.0 [110/2] via 130.50.50.1, 00:00:09, Vlan50
    10.0.0.0/8 is variably subnetted, 6 subnets, 2 masks
O IA   10.44.46.1/32 [110/3] via 130.50.50.1, 00:00:09, Vlan50
O IA   10.44.45.1/32 [110/3] via 130.50.50.1, 00:00:09, Vlan50
O IA   10.44.44.1/32 [110/3] via 130.50.50.1, 00:00:09, Vlan50
SW3#

SW4# show ip route ospf
    130.50.0.0/30 is subnetted, 1 subnets
O IA   130.50.50.0 [110/2] via 130.100.100.1, 00:00:07, Vlan100
    10.0.0.0/8 is variably subnetted, 6 subnets, 2 masks
O IA   10.33.34.1/32 [110/3] via 130.100.100.1, 00:00:07, Vlan100
O IA   10.33.35.1/32 [110/3] via 130.100.100.1, 00:00:07, Vlan100
O IA   10.33.33.1/32 [110/3] via 130.100.100.1, 00:00:07, Vlan100
```

Section 6: MPLS (7 Points)

- Leak network 10.1.1.0/24 from Sw1 VRF BLUE on PE R1 into the VRF RED on PE1; similarly, leak 10.44.44.0/24 from VRF RED into VRF BLUE on R6. Both Switch 1 and Switch 4 should receive the following routes:

```
SW1# show ip route | include 10.44.44.0
D EX  10.44.44.0/24 [170/XXXXXX] via 10.10.10.1, 00:00:27, Vlan10
SW1#
```

```
SW4# show ip route | include 10.1.1.0
```

```
O E1 10.1.1.0/24 [110/XX] via 130.100.100.1, 00:03:04, Vlan100
```

```
SW4#
```

Verify your configuration by pinging from VRF RED Sw4 10.44.44.1 to VRF BLUE Sw1 10.1.1.1 sw1. (5 points)

This is a straightforward VRF Export question with a slight twist for the attentive in that the OSPF route 10.44.44.0/24 originates from a Loopback interface on Switch4, so OSPF must be manipulated to treat this interface as a point-to-point network to advertise the /24 mask. The route-leaking is achieved by creation of export maps on the PE Routers R1 and R6, permitting the required routes from each VRF to the existing BLUE and RED VRF advertisements by adding them to the appropriate Route Target (RT) within MP-BGP by use of the `set extcommunity rt XXXXX:XXX additive` command. Example 3-20 details the required configuration on PE Routers R1, R6, and CE device Sw4; the resulting verification of the route advertisements and testing are also shown. If you have configured this correctly, as shown in Example 3-20, you have scored 5 points.

EXAMPLE 3-20 Selective VRF Export Configuration and Verification

```
Sw4(config)# interface Loopback0
Sw4(config-if)# ip ospf network point-to-point

R1(config)# ip vrf BLUE
R1(config-vrf)# export map SW1
R1(config-vrf)# access-list 10 permit 10.1.1.0 0.0.0.255
R1(config-vrf)# exit
R1(config)# route-map SW1 permit 10
R1(config-route-map)# match ip address 10
R1(config-route-map)# set extcommunity rt 65001:200 additive

R6(config)# ip vrf RED
R6(config-vrf)# export map SW4
R6(config-vrf)# access-list 10 permit 10.44.44.0 0.0.0.255
R6(config-vrf)# exit
R6(config)# route-map SW4 permit 10
R6(config-route-map)# match ip address 10
R6(config-route-map)# set extcommunity rt 65001:100 additive
```

! R1 is now sending 10.1.1.0 into VRF RED and R6 10.44.44.0 into VRF BLUE

R1# **show ip bgp vpnv4 vrf RED**

BGP table version is 33, local router ID is 120.100.1.1

Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
r RIB-failure, S Stale

Origin codes: i - IGP, e - EGP, ? - incomplete

Network	Next Hop	Metric	LocPrf	Weight	Path
Route Distinguisher: 65001:200 (default for vrf RED)					
*> 10.33.33.1/32	130.50.50.2	2		32768	?
*> 10.33.34.1/32	130.50.50.2	2		32768	?
*> 10.33.35.1/32	130.50.50.2	2		32768	?
*>i10.44.44.1/32	120.100.6.1	2	100		0 ?
*>i10.44.45.1/32	120.100.6.1	2	100		0 ?
*>i10.44.46.1/32	120.100.6.1	2	100		0 ?
*> 130.50.50.0/30	0.0.0.0		0		32768 ?
*>i130.100.100.0/30	120.100.6.1	0	100		0 ?

! No sign of the 10.1.1.0 route, clear the BGP session to kick start the export map

R1# **clear ip bgp ***

R1# **show ip bgp vpnv4 vrf RED**

BGP table version is 34, local router ID is 120.100.1.1

Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
r RIB-failure, S Stale

Origin codes: i - IGP, e - EGP, ? - incomplete

Network	Next Hop	Metric	LocPrf	Weight	Path
Route Distinguisher: 65001:200 (default for vrf RED)					
*> 10.1.1.0/24	10.10.10.2	50		32768	?
*> 10.33.33.1/32	130.50.50.2	2		32768	?
*> 10.33.34.1/32	130.50.50.2	2		32768	?
*> 10.33.35.1/32	130.50.50.2	2		32768	?
*>i10.44.44.1/32	120.100.6.1	2	100		0 ?
*>i10.44.45.1/32	120.100.6.1	2	100		0 ?
*>i10.44.46.1/32	120.100.6.1	2	100		0 ?
*> 130.50.50.0/30	0.0.0.0		0		32768 ?
*>i130.100.100.0/30	120.100.6.1	0	100		0 ?

R6# **show ip bgp vpnv4 vrf BLUE**

BGP table version is 35, local router ID is 120.100.6.1

Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
r RIB-failure, S Stale

Origin codes: i - IGP, e - EGP, ? - incomplete

Network	Next Hop	Metric	LocPrf	Weight	Path
Route Distinguisher: 65001:100 (default for vrf BLUE)					
*>i10.1.1.0/24	120.100.1.1	50	100	0	?
*>i10.1.2.0/24	120.100.1.1	50	100	0	?
*>i10.1.3.0/24	120.100.1.1	50	100	0	?
*> 10.2.2.0/24	10.10.20.2	50		32768	?
*> 10.2.3.0/24	10.10.20.2	50		32768	?
*> 10.2.4.0/24	10.10.20.2	50		32768	?
*>i10.10.10.0/30	120.100.1.1	0	100	0	?
*> 10.10.20.0/30	0.0.0.0		0	32768	?
*> 10.44.44.1/32	130.100.100.2	2		32768	?

! Notice the 10.44.44.0 route is actually listed as a host route, change the Loopback interface on Sw4 to a point-to-point for OSPF to advertise it correctly

SW4(config)# **interface lo0**

SW4(config-if)# **ip ospf network point-to-point**

R6# **show ip bgp vpnv4 vrf BLUE | include 10.44.44.0**

```
*> 10.44.44.0/24 130.100.100.2 2 32768 ?
```

Switch1# **show ip route | include 10.44.44.0**

```
D EX 10.44.44.0/24 [170/281856] via 10.10.10.1, 00:00:51, Vlan10
```

Switch1#

SW4# **show ip route | include 10.1.1.0**

```
O E1 10.1.1.0/24 [110/51] via 130.100.100.1, 00:02:45, Vlan100
```

! Now test with an extended ping to ensure the Loopback interface is used as the source

SW1# **ping**

Protocol [ip]:

Target IP address: **10.44.44.1**

Repeat count [5]:

```

Datagram size [100]:
Timeout in seconds [2]:
Extended commands [n]: y
Source address or interface: 10.1.1.1
Type of service [0]:
Set DF bit in IP header? [no]:
Validate reply data? [no]:
Data pattern [0xABCD]:
Loose, Strict, Record, Timestamp, Verbose[none]:
Sweep range of sizes [n]:
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.44.44.1, timeout is 2 seconds:
Packet sent with a source address of 10.1.1.1
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 8/9/12 ms

```

R1# show mpls forwarding-table vrf BLUE

Local tag	Outgoing tag or VC	Prefix or Tunnel Id	Bytes switched	tag	Outgoing interface	Next Hop
34	Untagged	10.1.3.0/24[V]	0		Gi0/0.10	10.10.10.2
35	Untagged	10.1.2.0/24[V]	0		Gi0/0.10	10.10.10.2
36	Aggregate	10.10.10.0/30[V]	0			
37	Untagged	10.1.1.0/24[V]	590		Gi0/0.10	10.10.10.2

R1# show mpls forwarding-table vrf RED

Local tag	Outgoing tag or VC	Prefix or Tunnel Id	Bytes switched	tag	Outgoing interface	Next Hop
38	Aggregate	130.50.50.0/30[V]	0			
39	Untagged	10.33.34.1/32[V]	0		Gi0/0.50	130.50.50.2
40	Untagged	10.33.35.1/32[V]	0		Gi0/0.50	130.50.50.2
41	Untagged	10.33.33.1/32[V]	0		Gi0/0.50	130.50.50.2

! Note the Routes are not leaked within the MPLS forwarding-table

R6# show mpls forwarding-table vrf BLUE

Local tag	Outgoing tag or VC	Prefix or Tunnel Id	Bytes switched	tag	Outgoing interface	Next Hop
34	Untagged	10.2.2.0/24[V]	0		Gi0/1.20	10.10.20.2
35	Untagged	10.2.3.0/24[V]	0		Gi0/1.20	10.10.20.2
36	Untagged	10.2.4.0/24[V]	0		Gi0/1.20	10.10.20.2

```

37      Aggregate    10.10.20.0/30[V]    0

R6# show mpls forwarding-table vrf RED
Local   Outgoing   Prefix           Bytes tag   Outgoing   Next Hop
tag     tag or VC   or Tunnel Id     switched   interface
38      Aggregate    130.100.100.0/30[V]    0
39      Untagged     10.44.46.1/32[V]      0      Gi0/1.100   130.100.100.2
40      Untagged     10.44.45.1/32[V]      0      Gi0/1.100   130.100.100.2
42      Untagged     10.44.44.0/24[V]     1534    Gi0/1.100   130.100.100.2

! Note the Routes are not leaked within the MPLS forwarding-table

```

- Configure your PE Routers R1 and R6 to ensure that the MPLS P routers are not listed as intermediate hops when a trace route is performed on your CE devices. (2 points)

By default, the MPLS network will be shown when a traceroute is performed. This can be changed, so only PE routers are shown as next hops with the **no mpls ip propagate-ttl** global command within your PE routers. Example 3-21 shows the default behavior and modified behavior after configuration from a trace route command issued on CE device SW1. If you have configured this correctly, as shown in Example 3-21, you have scored 2 points.

EXAMPLE 3-21 MPLS Traceroute Configuration and Testing

```

SW1# traceroute 10.2.2.1

Type escape sequence to abort.
Tracing the route to 10.2.2.1

 1 10.10.10.1 0 msec 0 msec 0 msec
 2 120.100.123.2 12 msec 12 msec 16 msec
 3 120.100.25.5 8 msec 12 msec 8 msec
 4 10.10.20.1 8 msec 8 msec 8 msec
 5 10.10.20.2 8 msec * 4 msec

R1(config)# no mpls ip propagate-ttl

R6(config)# no mpls ip propagate-ttl

SW1# traceroute 10.2.2.1

Type escape sequence to abort.

```

```
Tracing the route to 10.2.2.1

 1 10.10.10.1  4 msec  0 msec  0 msec
 2 10.10.20.1 12 msec  8 msec 12 msec
 3 10.10.20.2  4 msec  *   4 msec
```

Section 7: VPLS Simulation (10 Points)

- Switches 3 and 4 will have been configured to belong to the subnet of 1.1.1.0/24 in a previous question. Create an Xconnect attachment circuit on your PE Routers R1 and R6 for your CE devices (Sw3 Fe 0/19 1.1.1.1/24 and Sw4 Fe 0/19 1.1.1.2/24) to communicate using a secure Layer 2 tunneling solution (use version 3) across your Layer 3 network. You should use existing Loopback interfaces on your PE routers for peering over your MPLS network. Use a class template that configures a cookie size of 8 and a password of cisco, which will be used by a pseudowire class which Xconnects your required interfaces on your PE Routers R1 and R6. Be aware that the Sw3 resides in VLAN200 and Sw4 resides in VLAN400 in respective PE router subinterfaces. (10 points)

This question simulates VPLS and requires that L2TPv3 (Layer 2 Tunneling Protocol v3) is configured between your PE routers connecting the two subinterfaces that connect to Sw3 and Sw4 interfaces via Sw1 and Sw4 (VLAN200 and VLAN400, respectively). You might have considered using a VPLS-type solution, but the question dictates a secure Layer 2 tunneling solution and also provides you with an Xconnect clue and a version number. As such, it can only be L2TPv3. Sw3 and Sw4 will use a pseudowire to communicate over the IP network and logically will connect in the same Layer 2 domain. The PE routers have as directed a **l2tp-class** named SECURE. This configures the password to cisco and cookie size to 8; this class calls the pseudowire class **PW-CLASS**, which configures the encapsulation to **l2tpv3** in secure mode and sets the Loopback interfaces of the PE routers to be used for peering. The **xconnect** subinterface command binds the local PE interface to the remote PE Loopback with a vc-id (virtual channel ID), which in the exam-ple matches the subinterface number of the specific PE router. (You could have used any ID here.) It should be noted that Cisco Express Forwarding (CEF) must be enabled for the L2TPv3 feature to function correctly. Example 3-22 details the required PE configuration on Routers R1 and R2.

EXAMPLE 3-22 PE L2TPv3 Configuration

```
R1(config)# l2tp-class SECURE
R1(config-l2tp-class)# password 0 cisco
R1(config-l2tp-class)# cookie size 8
R1(config-l2tp-class)# pseudowire-class PW-CLASS
R1(config-pw-class)# encapsulation l2tpv3
R1(config-pw-class)# protocol l2tpv3 SECURE
```

```

R1(config-pw-class)# ip local interface Loopback0
R1(config-pw-class)# interface GigabitEthernet0/0.200
R1(config-subif)# xconnect 120.100.6.1 200 pw-class PW-CLASS

R6(config)# l2tp-class SECURE
R6(config-l2tp-class)# password 0 cisco
R6(config-l2tp-class)# cookie size 8
R6(config-l2tp-class)# pseudowire-class PW-CLASS
R6(config-pw-class)# encapsulation l2tpv3
R6(config-pw-class)# protocol l2tpv3 SECURE
R6(config-pw-class)# ip local interface Loopback0
R6(config-pw-class)# interface GigabitEthernet0/1.400
R6(config-subif)# xconnect 120.100.1.1 200 pw-class PW-CLASS

```

Example 3-23 shows the successful L2TPv3 session established between PE R1 to PE R6, yet the ping test from Sw3 to 1.1.1.2 fails. As the session is up, you can safely assume that there is a connectivity type issue between either Sw3 and PE R1 or Sw4 and PE R6, or possibly between both connections. The question does bring your attention to the fact that both CE devices reside in different VLANs, so this should give you a starting point in your investigation. When logging is enabled on Sw1 and Sw2 (these CE devices bring Sw3 and Sw4 FastEthernet 0/19 interfaces into VLAN200 and VLAN400, respectively), you can see spanning-tree inconsistencies exist between VLAN200 being “bridged” to VLAN400 via your L2TPv3 solution. Closer inspection reveals that spanning tree has actually blocked ports on Sw1 and Sw2 from PE Routers R1 and R6, respectively, even though you have previously allowed the local VLAN 200 and 400 through the trunk on PE Routers R1 and R6, respectively. The problem is actually resolved by enabling BPDU filtering on Sw1 with the **spanning-tree bpdupfilter enable** command on the trunk interface toward the PE Router R1. Enabling BPDU filtering on an interface is equivalent to disabling the spanning tree on an interface; it is possible to create bridging loops if this command is not correctly used. If you have configured this correctly, per Examples 3-22 and 3-23, you have scored 10 points.

EXAMPLE 3-23 PE and CE L2TPv3 Verification Testing and Configuration

```

R1# show l2t session

L2TP Session Information Total tunnels 1 sessions 1

LocID RemID Remote Name State Remote Address Port Sessions L2TP Class/
VPDN Group
51446 36190 R6 est 120.100.6.1 0 1 SECURE

```

```

LocID      RemID      TunID      Username, Intf/          State  Last Chg Uniq ID
                               Vcid, Circuit
51003      9619        51446      200, Gi0/0.200:200  est   00:24:40 1
    
```

R6# **show l2t session**

L2TP Tunnel and Session Information Total tunnels 1 sessions 1

```

LocID RemID Remote Name  State  Remote Address  Port  Sessions L2TP Class/
                               VPDN Group
36190 51446  R1                est   120.100.1.1      0      1      SECURE
    
```

```

LocID      RemID      TunID      Username, Intf/          State  Last Chg Uniq ID
                               Vcid, Circuit
9619       51003      36190      200, Gi0/1.400:400  est   00:25:26 1
    
```

SW3# **ping 1.1.1.2**

Type escape sequence to abort.
 Sending 5, 100-byte ICMP Echos to 1.1.1.2, timeout is 2 seconds:

!Make sure you are logging on your CE devices

SW1(config)# **logging console**

```

SW1#
03:22:19: %SPANTREE-2-RECV_PVID_ERR: Received BPDU with inconsistent peer vlan id 400 on FastEthernet0/1
VLAN200.
03:22:19: %SPANTREE-2-BLOCK_PVID_LOCAL: Blocking FastEthernet0/1 on VLAN0200. Inconsistent local vlan.
    
```

SW1# **show spanning-tree blockedports**

```

Name                Blocked Interfaces List
-----
VLAN0200            Fa0/1
    
```

Number of blocked ports (segments) in the system : 1

```
SW2#03:22:21: %SPANTREE-2-RECV_PVID_ERR: Received BPDU with inconsistent peer vlan id 200 on FastEthernet0/6
VLAN400.
```

```
03:22:21: %SPANTREE-2-BLOCK_PVID_PEER: Blocking FastEthernet0/6 on VLAN0200. Inconsistent peer vlan.
```

```
SW2# show spanning-tree blockedports
```

```
Name                Blocked Interfaces List
-----
VLAN0200            Fa0/6
VLAN0400            Fa0/6
```

```
Number of blocked ports (segments) in the system : 2
```

```
SW3# ping 1.1.1.2
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 1.1.1.2, timeout is 2 seconds:
```

```
.....
```

```
Success rate is 0 percent (0/5)
```

```
SW1# show spanning-tree blockedports
```

```
Name                Blocked Interfaces List
-----
VLAN0200            Fa0/1
```

```
Number of blocked ports (segments) in the system : 1
```

```
SW1(config)# int fast 0/1
```

```
SW1(config-if)# spanning-tree bpdupfilter enable
```

```
SW1(config-if)#03:33:57: %SPANTREE-2-UNBLOCK_CONSIST_PORT: Unblocking FastEthernet0/1 on VLAN0200. Port con-
sistency restored.
```

```
SW3# ping 1.1.1.2
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 1.1.1.2, timeout is 2 seconds:
```

```
.!!!!
```

```
Success rate is 80 percent (4/5), round-trip min/avg/max = 8/12/17 ms
```

Section 8: Multicast (10 Points)

- Configure your MPLS network for multicast support of the RED VRF using PIM sparse mode. PE Routers R1 and R6 should be configured to tunnel multicast traffic using an MDT address of 232.0.0.11 from CE device Switch 3 VLAN50 to CE device Sw4 VLAN100 over the RED VRF. Switch 4 should be configured to reply to an ICMP ping on its VLAN100 interface directed to 226.2.2.2 from Switch 3 VLAN50. It can be assumed that the mVRF bandwidth requirement is low; configure MDT appropriately. Ensure that PE Router R6's associated VLAN100 IP address is used as the rendezvous point (RP) for the RED VRF multicast traffic. (10 points)

Multicast support for MPLS VPNs is provided by configuring multicast routing within the core network. As directed, PIM sparse mode is required in your solution and should be enabled on all P router MPLS interfaces and P facing PE router MPLS interfaces. PIM sparse mode is also configured on the CE interfaces on VLAN50 and VLAN100 on Switches 3 and 4, respectively, and corresponding PE terminating interfaces on the PE Routers R1 and R6. PIM sparse mode is finally configured on the loopback interfaces of the PE Routers R1 and R6 as Multicast Distribution Tree (MDT) will tunnel between these interfaces.

Don't forget that multicast routing is enabled on the CE switches with the command **ip multicast-routing distributed** and on the routers with **ip multicast-routing**. The **mdt default group-address** is configured to 232.0.0.11 on PE Routers R1 and R6 within the RED VRF. Source Specific Multicast (SSM) is enabled on all MPLS routers with the command **ip pim ssm default** to allow transport of multicast information between all P and PE routers.

The question states that the mVRF (Multicast VRF) bandwidth requirement is low, which simply means that a Data MDT is not required in this solution. (These are used for high-bandwidth sources and limit the traffic received to the routers' part of the multicast tree.) You should also realize that a Data MDT is not required because there was no mention of threshold values or access-lists within the question, which are required for Data MDT configurations.

The address of 130.100.100.1 (R6 VRF RED) is used as the RP for the mVRF, and this is configured on both CE (Switch3 and Switch4) devices and both PE routers (R1 and R6) within the RED VRF.

CE device Switch 4 is finally configured with **ip igmp join-group 226.2.2.2** under its VLAN 100 interface for it to reply to a multicast ping from CE device Switch 3 over the MPLS VPN.

The question is comprehensive in the amount of items that require configuration, and it would be an easy mistake to miss tasks such as enabling PIM on the PE Loopback interfaces, where you might not immediately assume it is required. As with all questions, testing is key. Example 3-24 details the required configuration for the solution.

EXAMPLE 3-24 Multicast Configuration

```
! Initial Multicast Setup for the MPLS Core Routers
```

```
R1(config)# ip multicast-routing
R1(config-vrf)# interface Loopback0
R1(config-if)# ip pim sparse-mode
R1(config-if)# interface Serial0/0/0
R1(config-if)# ip pim sparse-mode

R2(config)# ip multicast-routing
R2(config)# interface s0/0
R2(config-if)# ip pim sparse-mode
R2(config-if)# interface s0/1
R2(config-if)# ip pim sparse-mode

R3(config)# ip multicast-routing
R3(config)# interface s0/0/0
R3(config-if)# ip pim sparse-mode
R3(config-if)# interface s0/0/1
R3(config-if)# ip pim sparse-mode

R4(config)# ip multicast-routing
R4(config)# interface gig0/0
R4(config-if)# ip pim sparse-mode
R4(config-if)# interface s0/0/1
R4(config-if)# ip pim sparse-mode

R5(config)# ip multicast-routing
R5(config)# interface gig0/0
R5(config-if)# ip pim sparse-mode
R5(config-if)# interface s0/0/1
R5(config-if)# ip pim sparse-mode

R6(config)# ip multicast-routing
R6(config)# interface Loopback0
R6(config-if)# ip pim sparse-mode
R6(config)# interface GigabitEthernet0/0
```

```
R6(config-if)# ip pim sparse-mode

! PE Specific mVRF and MDT Configuration

R1(config)# ip multicast-routing vrf RED
R1(config)# ip vrf RED
R1(config-vrf)# mdt default 232.0.0.11
R1(config-vrf)# interface GigabitEthernet0/0.50
R1(config-subif)# ip pim sparse-mode
R1(config-subif)# exit
R1(config)# ip pim vrf RED rp-address 130.100.100.1
R1(config)# ip pim ssm default

R6(config)# ip vrf RED
R6(config-vrf)# mdt default 232.0.0.11
R6(config-vrf)# interface GigabitEthernet0/1.100
R6(config-subif)# ip pim sparse-mode
R6(config-subif)# exit
R6(config)# ip pim vrf RED rp-address 130.100.100.1
R6(config)# ip pim ssm default

! CE Specific Configuration

SW3(config)# ip multicast-routing distributed
SW3(config)# int vlan 50
SW3(config-if)# ip pim sparse-mode
SW3(config-if)# exit
SW3(config)# ip pim rp-address 130.100.100.1

SW4(config)# ip multicast-routing distributed
SW4(config)# interface vlan 100
SW4(config-if)# ip pim sparse-mode
SW4(config-if)# ip igmp join-group 226.2.2.2
SW4(config-if)# exit
SW4(config)# ip pim rp-address 130.100.100.1
```

Example 3-25 details the testing for the solution; the MDT tunnel is detailed and shown as an interface used for PIM adjacency between the PE routers. If you have configured your solution per Example 3-25 and can successfully ping between Switch 3 and Switch 4, you have scored 10 points.

EXAMPLE 3-25 Multicast Testing

```
R6# show ip pim vrf RED neigh
PIM Neighbor Table
Mode: B - Bidir Capable, DR - Designated Router, N - Default DR Priority,
      S - State Refresh Capable
Neighbor          Interface                Uptime/Expires    Ver  DR
Address
Prio/Mode
130.100.100.2     GigabitEthernet0/1.100    00:02:08/00:01:34 v2   1 / DR S
120.100.1.1       Tunnel1                    00:00:05/00:01:39 v2   1 / S

R1# ping vrf RED 226.2.2.2

Type escape sequence to abort.
Sending 1, 100-byte ICMP Echos to 226.2.2.2, timeout is 2 seconds:

Reply to request 0 from 130.100.100.2, 12 ms

SW3# ping 226.2.2.2

Type escape sequence to abort.
Sending 1, 100-byte ICMP Echos to 226.2.2.2, timeout is 2 seconds:

Reply to request 0 from 130.100.100.2, 9 ms
SW3# show ip pim rp
Group: 226.2.2.2, RP: 130.100.100.1, v2, uptime 00:00:37, expires never
Group: 224.0.1.40, RP: 130.100.100.1, v2, uptime 01:01:24, expires never

R1# show ip pim mdt bgp
Peer (Route Distinguisher + IPv4)                Next Hop
MDT group 232.0.0.11
2:65001:200:120.100.6.1                            120.100.6.1

R6# show ip pim mdt bgp
Peer (Route Distinguisher + IPv4)                Next Hop
```

```
MDT group 232.0.0.11
```

```
2:65001:200:120.100.1.1
```

```
120.100.1.1
```

Section 9: IPv6 (6 Points)

- Configure the following IPv6 address on the PE Routers R1 and R6, and implement IPv6 over MPLS between the 6PE routers to advertise the prefixes between 6PEs. Ensure your Loopback IPv6 addresses are used to source any locally generated IPv6 traffic. (6 points)

```
R1 Lo0 2010:C15:C0:1::1/64
```

```
R1 Gi0/0.10 2010:C15:C0:11::1/64
```

```
R6 Lo0 2010:C15:C0:6::1/64
```

```
R6 Gi1/0.20 2010:C15:C0:62::1/64
```

A relatively straightforward IPv6 question, there is no IPv6 redistribution or complex issues to deal with. The question directs you to configure IPv6 onto your VRF BLUE interfaces of the PE routers. You would usually extend this IPv6 domain into your CE devices, but the switches in this lab cannot run IPv6. IPv6 over MPLS backbones enables isolated IPv6 domains to communicate with each other over an MPLS IPv4 core network. To ensure the Loopback IPv6 addresses of the PE routers are used to source locally generated IPv6 traffic, the PE routers are configured with **mpls ipv6 source-interface Loopback0**. MP-BGP is used to advertise the IPv6 prefixes between PE routers, and the configuration is virtually identical to that of IPv4. Aggregate label binding and advertisement is enabled for IPv6 prefixes using the **neighbor send-label** command. Connected IPV6 routes are redistributed using BGP with the **network** command under the IPv6 **address-family**, and IPv6 routing and IPv6 cef must be enabled on your PE routers. If you have configured your routers correctly, per Example 3-26, you have scored 6 points.

EXAMPLE 3-26 PE IPv6 Configuration and Verification

```
R1(config)# ipv6 unicast-routing
R1(config)# ipv6 cef
R1(config)# mpls ipv6 source-interface Loopback0
R1(config)# interface loopback0
R1(config-if)# ipv6 add 2010:C15:C0:1::1/64
R1(config-if)# interface GigabitEthernet0/0.10
R1(config-subif)# ipv6 address 2010:C15:C0:11::1/64
R1(config-subif)# router bgp 65001
R1(config-router)# no bgp default ipv4-unicast
```

```
R1(config-router)# address-family ipv6
R1(config-router-af)# neighbor 120.100.6.1 activate
R1(config-router-af)# neighbor 120.100.6.1 send-label
R1(config-router-af)# network 2010:C15:C0:11::0/64
R1(config-router-af)# network 2010:C15:C0:1::/64
R1(config-router-af)# exit-address-family
```

```
R6(config)# ipv6 unicast-routing
R6(config)# ipv6 cef
R6(config)# mpls ipv6 source-interface Loopback0
R6(config)# interface loopback0
R6(config-if)# ipv6 add 2010:C15:C0:6::1/64
R6(config-if)# interface GigabitEthernet1/0.20
R6(config-subif)# ipv6 address 2010:C15:C0:62::1/64
R6(config-subif)# router bgp 65001
R6(config-router)# no bgp default ipv4-unicast
R6(config-router)# address-family ipv6
R6(config-router-af)# neighbor 120.100.1.1 activate
R6(config-router-af)# neighbor 120.100.1.1 send-label
R6(config-router-af)# network 2010:C15:C0:62::/64
R6(config-router-af)# network 2010:C15:C0:6::/64
R6(config-router-af)# exit-address-family
```

```
R1# show ip bgp ipv6 unicast
```

```
BGP table version is 5, local router ID is 120.100.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
```

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 2010:C15:C0:1::/64	::	0		32768	i
*>i2010:C15:C0:6::/64	::FFFF:120.100.6.1	0	100	0	i
*> 2010:C15:C0:11::/64	::	0		32768	i
*>i2010:C15:C0:62::/64	::FFFF:120.100.6.1				

```

0 100 0 i

R6# show ip bgp ipv6 unicast
BGP table version is 5, local router ID is 120.100.6.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
*>i2010:C15:C0:1::/64
                   ::FFFF:120.100.1.1
                   0 100 0 i
*> 2010:C15:C0:6::/64
                   ::
                   0 32768 i
*>i2010:C15:C0:11::/64
                   1.1
                   0 100 0 i
                   ::
                   0 32768 i

```

R1# ping ipv6 2010:C15:C0:62::1

```

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2010:C15:C0:62::1, timeout is 2 seconds:
!!!!

```

R1# ping ipv6 ::C0:6::1

```

Success rate is 100 percent (5/5), round-trip min/avg/max = 8/8/12 ms

```

```

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2010:C15:C0:6::1, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 8/8/12 ms

```

R6# ping ipv6 2010:C15:C0:11::1

```

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2010:C15:C0:11::1, timeout is 2 seconds:
!!!!

```

```

Success rate is 100 percent (5/5), round-trip min/avg/max = 8/8/12 ms

```

R6# ping ipv6 2010:C15:C0:1::1


```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2010:C15:C0:1::1, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 8/9/12 ms
```

R1# show ipv6 route

```
IPv6 Routing Table - 8 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B -
       BGP U - Per-user Static route
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
       D - EIGRP, EX - EIGRP external
C 2010:C15:C0:1::/64 [0/0]
   via ::, Loopback0
L 2010:C15:C0:1::1/128 [0/0]
   via ::, Loopback0
B 2010:C15:C0:6::/64 [200/0]
   via ::FFFF:120.100.6.1, IPv6-mpls
C 2010:C15:C0:11::/64 [0/0]
   via ::, GigabitEthernet0/0.10
L 2010:C15:C0:11::1/128 [0/0]
   via ::, GigabitEthernet0/0.10
B 2010:C15:C0:62::/64 [200/0]
   via ::FFFF:120.100.6.1, IPv6-mpls
L FE80::/10 [0/0]
   via ::, Null0
L FF00::/8 [0/0]
   via ::, Null0
```

R6# show ipv6 route

```
IPv6 Routing Table - 8 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B -
       BGP U - Per-user Static route
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
       D - EIGRP, EX - EIGRP external
B 2010:C15:C0:1::/64 [200/0]
```

```

    via ::FFFF:120.100.1.1, IPv6-mpls
C 2010:C15:C0:6::/64 [0/0]
    via ::, Loopback0
L 2010:C15:C0:6::1/128 [0/0]
    via ::, Loopback0
B 2010:C15:C0:11::/64 [200/0]
    via ::FFFF:120.100.1.1, IPv6-mpls
C 2010:C15:C0:62::/64 [0/0]
    via ::, GigabitEthernet0/1.10
L 2010:C15:C0:62::1/128 [0/0]
    via ::, GigabitEthernet0/1.20
L FE80::/10 [0/0]
    via ::, Null0
L FF00::/8 [0/0]
    via ::, Null0

```

Section 10: QoS (13 Points)

- Create the following QoS profile on your PE Router R1 for traffic egressing to your CE device connected to the BLUE VRF. Use an appropriate method of prioritizing DSCP traffic so that AF31 packets are statistically dropped more frequently than AF32 during congestion and reduce the effects of TCP global synchronization within your MISSION-CRITICAL and solely reduce the effect of TCP global synchronization within the DEFAULT class: (7 points)

Class	DSCP Value	% of Bandwidth Assigned
VOICE	EF, CS5	35
MISSION-CRITICAL	CS6, AF31, AF32, CS3	40
DEFAULT	Any	25

This is a 3 Class PE-to-CE QoS question that requires assigning traffic to queues based on DSCP values into the listed classes and assignment of bandwidth on a per-class basis. DSCP prioritization is achieved in the MISSION-CRITICAL class by enabling WRED with the **random-detect dscp-based** command, whereby lower-priority DSCP traffic will be dropped more aggressively than higher priority under congestion, thus reducing the effect of global synchronization. A similar non-DSCP-based effect is achieved within the DEFAULT class by use of the **random-detect** command. The policy-map is applied outbound on the PE interface connecting to the BLUE VRF CE device. Example 3-27 details the required configuration on PE Router R1. If you have configured this correctly, you have scored 7 points.

EXAMPLE 3-27 PE to CE QoS Configuration

```

R1(config)# class-map match-any VOICE
R1(config-cmap)# match ip dscp ef
R1(config-cmap)# match ip dscp cs5
R1(config-cmap)# class-map match-any MISSION-CRITICAL
R1(config-cmap)# match ip dscp cs6
R1(config-cmap)# match ip dscp af31
R1(config-cmap)# match ip dscp af32
R1(config-cmap)# match ip dscp cs3
R1(config-cmap)# policy-map PE-CE
R1(config-pmap)# class VOICE
R1(config-pmap-c)# priority percent 35
R1(config-pmap-c)# class MISSION-CRITICAL
R1(config-pmap-c)# bandwidth percent 40
R1(config-pmap-c)# random-detect dscp-based
R1(config-pmap-c)# class class-default
R1(config-pmap-c)# bandwidth percent 25
R1(config-pmap-c)# random-detect
R1(config-pmap-c)# exit
R1(config-pmap)# exit
R1(config)# interface GigabitEthernet0/0.10
R1(config-subif)# service-policy output CE-PE

```

- Create the following QoS profile on your PE Router R1 for traffic ingressing from your CE device connected to the BLUE VRF into the MPLS network. The total aggregate speed from the CE to PE should be restricted to 1 Mbps:

Class	CIR (bps)
VOICE	350,000
MISSION-CRITICAL	400,000
DEFAULT	250,000

Traffic in the VOICE class within the detailed CIR should have the MPLS EXP set to 5 and above discarded. Traffic in the MISSION-CRITICAL class within the detailed CIR should have the MPLS EXP set to 3 and above set to 7. Traffic in the DEFAULT class within the detailed CIR should have the MPLS EXP set to 0 and above set to 4. (6 points)

This is a DiffServ Tunneling question which requires that the classes you have configured in the previous question be policed to an aggregate of 1 Mbps and have their MPLS EXP values adjusted. The policy-map is applied to the input interface of the PE router, which connects to the BLUE VRF CE device and affects the traffic as it flows through the MPLS network. Example 3-28 details the required configuration on PE Router R1. If you have configured this correctly, you have scored 6 points.

EXAMPLE 3-28 CE to PE QoS Configuration

```
R1(config)# policy-map CE-PE-SHAPE
R1(config-pmap)# class VOICE
R1(config-pmap-c)# police cir 350000
R1(config-pmap-c-police)# conform-action set-mpls-exp-topmost-transmit 5
R1(config-pmap-c-police)# exceed-action drop
R1(config-pmap-c-police)# class MISSION-CRITICAL
R1(config-pmap-c)# police cir 400000
R1(config-pmap-c-police)# conform-action set-mpls-exp-topmost-transmit 3
R1(config-pmap-c-police)# exceed-action set-mpls-exp-topmost-transmit 7
R1(config-pmap-c-police)# class class-default
R1(config-pmap-c)# police cir 250000
R1(config-pmap-c-police)# conform-action set-mpls-exp-topmost-transmit 0
R1(config-pmap-c-police)# exceed-action set-mpls-exp-topmost-transmit 4
R1(config-pmap-c-police)# interface GigabitEthernet0/0.10
R1(config-subif)# service-policy input CE-PE-SHAPE
```

Section 11: Security (13 Points)

- Create three new Loopback IP addresses of loopback1 on R4, R5, and R6—use IP addresses of 4.4.4.4/24, 5.5.5.5/24, and 6.6.6.6/24, respectively. Use EIGRP to advertise the loopback networks between routers over a common GRE tunnel network of 100.100.100.X / 24 (X = router number) sourced from each router's common Ethernet interface using IPsec to encrypt all traffic between the loopback networks using a preshared isakmp key of CCIE. Use an IPsec transform-set of esp-des esp-md5-hmac on each router. R6 needs to be a hub router, with R4 and R5 effectively being spoke routers in your solution. You are not permitted to enable EIGRP on your Ethernet interfaces between routers. Spoke routers must be able to communicate with each other directly using dynamic IPsec connections with the aid of NHRP at the hub, whereas hub-to-spoke IPsec connections should be permanent. The hub router should provide all necessary direct next-hop information to the spoke routers when they are required to communicate between themselves. NHRP should be authenticated with a password of

SECRET. Use an MTU of 1416 for your secure traffic, an NHRP timeout of 100 seconds for spoke replies, and a delay of 2mS on the tunnel network. Test your solution by extended pings sourced from the configured Loopback interfaces. (10 points)

This is a classic Dynamic Multipoint VPN (DMVPN) question in which a hub-and-spoke design is used with Next Hop Resolution Protocol (NHRP) for the spoke routers to communicate with each other. You have numerous tasks to perform, so this could be the kind of question that is best saved until later and tackled if you have time. The question dictates that you configure a tunnel network 100.100.100.0/24 in which to advertise each router's new Loopback network over GRE and EIGRP sourced from the common Ethernet interfaces, which is uncomplicated; the complexity begins when you enable IPsec and NHRP. The **crypto isakmp policy** command configures the preshared key to **CCIE** and sets the transform-set with the required parameters of **esp-des esp-md5-hmac**, which are applied to the tunnel interface by the use of the **tunnel protection ipsec profile IPSEC** command. The MTU is fixed at 1416 as directed within the question on the tunnel interfaces to allow for overhead of the VPN connection.

A delay of 2000 is configured on each tunnel interface as directed in the question, which is 2mS, so be aware of the unit values, which are micro seconds. The tunnel source of each router is the common Ethernet network 120.100.45. Because the spoke routers will terminate their connection to the hub on the same interface, the tunnel mode must be set to **tunnel mode gre multipoint**. NHRP is enabled on the tunnel interface of each router with an identical network ID to match the broadcast domain for all three routers, and the authentication password is set to **SECRET** as directed within the question. The command **ip nhrp map multicast dynamic** permits the registration of the multicast address for EIGRP during boot up or initiation of spoke-to-hub sessions. The **ip nhrp holdtime 100** command sets the NHRP time for a spoke to keep the NHRP reply to 100 seconds and is configured on the hub-and-spoke routers.

The required configuration for the Loopback and tunnel interfaces and the DMVPN is detailed in Example 3-29.

EXAMPLE 3-29 DMVPN Configuration

```
R4(config)# interface loopback1
R4(config-if)# ip add 4.4.4.4 255.255.255.0
R4(config-if)# router eigrp 1
R4(config-router)# no auto-summary
R4(config-router)# network 100.100.100.0 0.0.0.255
R4(config-router)# network 4.4.4.0 0.0.0.255

R5(config)# interface loopback1
R5(config-if)# ip address 5.5.5.5 255.255.255.0
R5(config-if)# router eigrp 1
```

```
R5(config-router)# no auto-summary
R5(config-router)# network 100.100.100.0 0.0.0.255
R5(config-router)# network 5.5.5.0 0.0.0.255

R6(config)# interface loopback1
R6(config-if)# ip address 6.6.6.6 255.255.255.0
R6(config-if)# router eigrp 1
R6(config-router)# no auto-summary
R6(config-router)# network 100.100.100.0 0.0.0.255
R6(config-router)# network 6.6.6.0 0.0.0.255

R6(config)# crypto isakmp policy 1
R6(config-isakmp)# authentication pre-share
R6(config-isakmp)# crypto isakmp key CCIE address 0.0.0.0
R6(config-isakmp)# crypto ipsec transform-set DMVPN esp-des esp-md5-hmac
R6(cfg-crypto-trans)# crypto ipsec profile IPSEC
R6(ipsec-profile)# set transform-set DMVPN
R6(ipsec-profile)# interface Tunnel0
R6(config-if)# ip address 100.100.100.6 255.255.255.0
R6(config-if)# ip mtu 1416
R6(config-if)# ip nhrp authentication SECRET
R6(config-if)# ip nhrp map multicast dynamic
R6(config-if)# ip nhrp network-id 10
R6(config-if)# ip nhrp holdtime 100
R6(config-if)# delay 2000
R6(config-if)# tunnel source gig 0/0
R6(config-if)# tunnel mode gre multipoint
R6(config-if)# tunnel key 1
R6(config-if)# tunnel protection ipsec profile IPSEC

R4(config)# crypto isakmp policy 1
R4(config-isakmp)# authentication pre-share
R4(config-isakmp)# crypto isakmp key CCIE address 0.0.0.0
R4(config-isakmp)# crypto ipsec transform-set DMVPN esp-des esp-md5-hmac
R4(cfg-crypto-trans)# crypto ipsec profile IPSEC
R4(ipsec-profile)# set transform-set DMVPN
R4(ipsec-profile)# interface Tunnel0
R4(config-if)# ip address 100.100.100.4 255.255.255.0
```

```
R4(config-if)# ip mtu 1416
R4(config-if)# ip nhrp authentication SECRET
R4(config-if)# ip nhrp map 100.100.100.6 120.100.45.6
R4(config-if)# ip nhrp map multicast 120.100.45.6
R4(config-if)# ip nhrp network-id 10
R4(config-if)# ip nhrp holdtime 100
R4(config-if)# ip nhrp nhs 100.100.100.6
R4(config-if)# delay 2000
R4(config-if)# tunnel source gig 0/0
R4(config-if)# tunnel mode gre multipoint
R4(config-if)# tunnel key 1
R4(config-if)# tunnel protection ipsec profile IPSEC

R5(config)# crypto isakmp policy 1
R5(config-isakmp)# authentication pre-share
R5(config-isakmp)# crypto isakmp key CCIE address 0.0.0.0
R5(config-isakmp)# crypto ipsec transform-set DMVPN esp-des esp-md5-hmac
R5(cfg-crypto-trans)# crypto ipsec profile IPSEC
R5(ipsec-profile)# set transform-set DMVPN
R5(ipsec-profile)# interface Tunnel0
R5(config-if)# ip address 100.100.100.5 255.255.255.0
R5(config-if)# ip mtu 1416
R5(config-if)# ip nhrp authentication SECRET
R5(config-if)# ip nhrp map 100.100.100.6 120.100.45.6
R5(config-if)# ip nhrp map multicast 120.100.45.6
R5(config-if)# ip nhrp network-id 10
R5(config-if)# ip nhrp holdtime 100
R5(config-if)# ip nhrp nhs 100.100.100.6
R5(config-if)# delay 2000
R5(config-if)# tunnel source gig 0/0
R5(config-if)# tunnel mode gre multipoint
R5(config-if)# tunnel key 1
R5(config-if)# tunnel protection ipsec profile IPSEC
```

Example 3-30 details the EIGRP routes received on all routers. As can be seen, the hub router shows both spoke networks, yet each spoke router discovers only the hub network; this is a classic split-horizon issue. The hub Router R6 must be configured to disable the split-horizon behavior to ensure the spoke routers receive each other's routes. How-

ever, the question dictates that spoke routers should be able to communicate “directly.” As shown in Example 3-30, the next hop for spoke networks show as the hub router 100.100.100.6 for each spoke network. The command **no ip next-hop-self eigrp 1** on the hub Router R6 ensures that the spoke routers are used as next hops when spoke-to-spoke communication is required, and this will enable the dynamic IPsec peering between spokes as directed in the question.

EXAMPLE 3-30 DMVPN Spoke-to-Spoke Routing

```
R4# show ip route eigrp
 6.0.0.0/24 is subnetted, 1 subnets
D       6.6.6.0 [90/285084416] via 100.100.100.6, 00:02:42, Tunnel0
```

```
R5# show ip route eigrp
 6.0.0.0/24 is subnetted, 1 subnets
D       6.6.6.0 [90/285084416] via 100.100.100.6, 00:00:50, Tunnel0
```

```
R6# show ip route eigrp
 4.0.0.0/24 is subnetted, 1 subnets
D       4.4.4.0 [90/285084416] via 100.100.100.4, 00:03:06, Tunnel0
 5.0.0.0/24 is subnetted, 1 subnets
D       5.5.5.0 [90/285084416] via 100.100.100.5, 00:01:02, Tunnel0
```

!R6 has both spoke routes yet each spoke (R4 and R5) only have the hub network route, !a classic split horizon issue.

```
R6(config)# interface tunnel0
R6(config-if)# no ip split-horizon eigrp 1
```

```
R4# show ip route eigrp
 5.0.0.0/24 is subnetted, 1 subnets
D       5.5.5.0 [90/285596416] via 100.100.100.6, 00:00:22, Tunnel0
 6.0.0.0/24 is subnetted, 1 subnets
D       6.6.6.0 [90/285084416] via 100.100.100.6, 00:04:14, Tunnel0
```

```
R5# show ip route eigrp
 4.0.0.0/24 is subnetted, 1 subnets
D       4.4.4.0 [90/285596416] via 100.100.100.6, 00:00:33, Tunnel0
 6.0.0.0/24 is subnetted, 1 subnets
D       6.6.6.0 [90/285084416] via 100.100.100.6, 00:02:20, Tunnel0
R5#
```

```
! The next-hop for spoke to spoke routes shows as the hub router (100.100.100.6) yet !the question
states traffic must flow directly between spokes so the next-hop must be !modified
```

```
R6(config)# interface tunnel 0
R6(config-if)# no ip next-hop-self eigrp 1

R4# show ip route eigrp
 5.0.0.0/24 is subnetted, 1 subnets
D       5.5.5.0 [90/285596416] via 100.100.100.5, 00:00:28, Tunnel0
 6.0.0.0/24 is subnetted, 1 subnets
D       6.6.6.0 [90/285084416] via 100.100.100.6, 00:00:29, Tunnel0

R5# show ip route eigrp
 4.0.0.0/24 is subnetted, 1 subnets
D       4.4.4.0 [90/285596416] via 100.100.100.4, 00:00:39, Tunnel0
 6.0.0.0/24 is subnetted, 1 subnets
D       6.6.6.0 [90/285084416] via 100.100.100.6, 00:00:39, Tunnel0
```

Example 3-31 shows the isakmp IPsec connection on spoke Router R5 to the hub. To bring up a dynamic isakmp IPsec connection to the other spoke Router R4, an extended ping is required from Loopback interface to Loopback interface.

This question was extremely complex and is the reason why it was weighted so heavily. You had multiple items to configure within the standard DMVPN solution, such as split-horizon. It should make you realize the importance of reading the question a number of times and taking the time to test your configurations to ensure you have successfully answered the question. If you have configured your routers correctly, as detailed in Examples 3-29 and 3-30, congratulations, and you have earned a hefty 10 points.

EXAMPLE 3-31 DMVPN Spoke-to-Spoke Testing

```
R5# show crypto map
Crypto Map "Tunnel0-head-0" 65536 ipsec-isakmp
  Profile name: IPSEC
  Security association lifetime: 4608000 kilobytes/3600 seconds
  PFS (Y/N): N
  Transform sets={
    DMVPN,
  }
```

```

Crypto Map "Tunnel0-head-0" 65537 ipsec-isakmp
  Map is a PROFILE INSTANCE.
  Peer = 120.100.45.6
  Extended IP access list
    access-list permit gre host 120.100.45.5 host 120.100.45.6
  Current peer: 120.100.45.6
  Security association lifetime: 4608000 kilobytes/3600 seconds
  PFS (Y/N): N
  Transform sets={
    DMVPN,
  }
  Interfaces using crypto map Tunnel0-head-0:
    Tunnel0

```

```

R5# show crypto isakmp sa
IPv4 Crypto ISAKMP SA

```

```

120.100.45.6 120.100.45.5 QM_IDLE 4001 0 ACTIVE
dst src state conn-id slot status

```

```

IPv6 Crypto ISAKMP SA

```

!R5 spoke router only has a connection to the Hub router. An extended ping sourced from the loopback interface of one spoke to another is required to bring up the dynamic spoke to spoke connection.

```

R5#ping

```

```

Protocol [ip]:
Target IP address: 4.4.4.4
Repeat count [5]:
Datagram size [100]:
Timeout in seconds [2]:
Extended commands [n]: y
Source address or interface: 5.5.5.5
Type of service [0]:
Set DF bit in IP header? [no]:
Validate reply data? [no]:
Data pattern [0xABCD]:
Loose, Strict, Record, Timestamp, Verbose[none]:
Sweep range of sizes [n]:
Type escape sequence to abort.

```



```

Sending 5, 100-byte ICMP Echos to 4.4.4.4, timeout is 2 seconds:
Packet sent with a source address of 5.5.5.5
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

```

```

R5# show crypto isakmp sa
IPv4 Crypto ISAKMP SA
dst          src          state          conn-id slot status
120.100.45.5 120.100.45.4 QM_IDLE        4002   0 ACTIVE
120.100.45.6 120.100.45.5 QM_IDLE        4001   0 ACTIVE

IPv6 Crypto ISAKMP SA

R5# show crypto isakmp sa
IPv4 Crypto ISAKMP SA
dst          src          state          conn-id slot status
120.100.45.5 120.100.45.4 QM_IDLE        4002   0 ACTIVE
120.100.45.6 120.100.45.5 QM_IDLE        4001   0 ACTIVE

IPv6 Crypto ISAKMP SA

```

- The network manager of your network cannot justify a full security implementation but wants to implement a solution that provides only a password prompt from R1 when the keyboard entry 1 is entered on the console port (as opposed to the normal CR/Enter key). Configure R1 appropriately. (3 points)

This question makes use of the **activation-character** command on the console port. This is a nasty question because the CLI entry requires an ASCII entry; you'd need to search to discover that ASCII numeric figures (0 to 9) are prefixed by the binary value of 0011, so a value of 1 (0001) would be 00110001; as such the decimal conversion is $32 + 16 + 1 = 49$. A good question to use the (?) on the CLI for clues and your documentation CD or search facility in the lab if you were not aware of this feature. If you have configured this correctly per Example 3-32, you have scored 3 points.

EXAMPLE 3-32 R1 Console Activation-Character Configuration

```

R1(config)# line con 0
R1(config-line)# activation-character ?
CHAR or <0-127> Activation character or its decimal equivalent
R1(config-line)# activation-character 49

```

Lab 3 Wrap-Up

So how did it go? Did you run out of time? Did you manage to finish but miss what was actually required? If you scored more than 80, well done. If you accomplished this within the time frame of 8 hours or less, you will be prepared for any scenario that you are likely to face during the 5 1/2 hours of the Configuration section of the actual exam. Remember that the Troubleshooting section on the v4.0 exam is a separate section to the configuration with a different scenario, and you will have 2 hours to complete this. This lab was designed to ensure you troubleshoot your own work as you progress through the questions.

Did you manage to configure items such as disabling split horizon for DMVPN and the area ID for OSPF? This attention to detail and complete understanding of the protocols will ultimately earn you your number.

Chapter 4 Summary

Are You Ready?

This became a well-known Cisco slogan that identified the Internet revolution. By the end of these practice exams, you should have a good idea of whether you are ready. Did you feel confident working through the questions, or was it a complete shock to the system? Are you more used to being spoon-fed solitary scenarios than actually having to analyze questions and piece together parts of a complex network jigsaw?

Life is full of challenges. During your education and career, the CCIE Certification is as tough as it gets. The exam is designed to test your technical skills, your understanding and analysis of complex topologies, and your capacity to build and troubleshoot a network with IP routing protocols and features. You need to achieve a minimum score of 80 percent to pass.

Further Reading

The following Cisco Press titles are on topics appearing on the CCIE exam blueprint. These books are not required study resources, but they can be used to build knowledge in certain areas.

CCIE Routing and Switching Exam Certification Guide, Fourth Edition

CCIE Routing and Switching Exam Quick Reference, Second Edition

CCIE Routing and Switching Troubleshooting Practice Labs

Routing TCP/IP, Volume I, 2/e Routing

TCP/IP, Volume II Troubleshooting IP

Routing Protocols

Inside Cisco IOS Software Architecture

Cisco LAN Switching

Cisco OSPF Command and Configuration Handbook

Cisco BGP-4 Command and Configuration Handbook

Cisco Router Configuration Handbook, Second Edition

Cisco LAN Switching Configuration Handbook, Second Edition

Developing IP Multicast Networks, Volume I

Internet Routing Architectures, Second Edition

MPLS and VPN Architectures

MPLS and VPN Architectures, Volume II

Cisco Catalyst QoS

End-to-End QoS Network Design

Deploying IPv6 Networks

Network Security Technologies and Solutions

Help and Advice

- Look at http://www.cisco.com/web/learning/le3/ccie/rs/lab_exam.html for the latest information regarding the CCIE Certification, which includes suggested training and reading.
- Keep your schedule flexible during your rack time. Include time for breaks and relaxation—you will often find that five minutes away from the keyboard can help you consider possible solutions. Most important, do not forget the people you care for and make time for them, too.
- Build your study plan based on a balance between theory and practice. You need to understand the concepts through the theory; then consolidate this during your rack time.

- Begin with simple topics in isolation; then work up to complex lab scenarios. Spend as much time repeating your configurations as possible to improve your speed and ability to perform basic configurations with your eyes shut. This will save you time for where you need it during the exam.
- Explore the Cisco CD documentation or the URL <http://www.cisco.com/univercd/home/home.htm>. This will be your research lifeline during the exam where you can find information, concepts, and samples regarding all technologies involved in the exam.
- Start to plan for your exam at least six months before the lab date.
- If you find these practice labs have highlighted weak areas, do not be afraid to postpone your lab date.

How Can I Schedule My CCIE Lab Exam?

Go to http://www.cisco.com/web/learning/le3/ccie/rs/lab_exam.html, and you can find all the information on how to schedule your exam including locations, start times, and more. You must have a CCO user ID, your CCIE written exam date, and score to be able to view your profile and schedule your exam.

The Day Before...

If you are traveling to take your exam, try to arrive the day before to familiarize yourself with the area. Take a tour to the lab location, so you won't be late on the day; the last thing you need is to arrive flustered. The day before is a day to be relaxed and not to attempt any last-minute studying. Have a light dinner and try to have a good night's sleep. Most important, save the beer until after the exam; pass or fail you will feel like one or two for sure. The CCIE exam might be the reason why Stella Artois is so popular in Brussels!

The Day of the Exam

On the day of the exam, you should plan to arrive at least 15 minutes before the exam begins for registration. The proctor will walk you to the lab and give you a briefing before the exam starts, telling you about the lab environment, on which rack or station you will be working, and the general guidelines for the day.

The proctor will not discuss solutions or possible solutions for a given question with you. The proctor will be available to help you understand the wording or meaning of the questions, make sure the backbone routers are working properly, and the hardware and software on your rack are working perfectly so your exam runs smoothly. Ask the proctor for any

assistance or verification; the worst he or she can say is, “Sorry, everything looks okay from my side; please check your configuration.” Read the entire exam before you start to get the bigger picture, ensuring you fully understand each question and its requirements. Begin by performing easier tasks, leaving the most difficult for later. Take some small breaks during the morning and the afternoon to refresh yourself and relieve the stress.

Pass or Fail, What Next?

If you pass, you certainly have something to celebrate; you have just joined a very elite club that will in no doubt enhance your career. You have achieved the highest level of certification in the networking world and should aim to continue your thirst for knowledge that sets you apart from your peers, but take a break before starting your next CCIE track!

If you failed, don't worry and don't take it personally; most people fail the first time around. You will have to put it down to experience and get back on the keyboard as soon as you can to work out what went wrong. You will more than likely be successful the next time and will ultimately become a better engineer for your extra rack time.

I hope these practice exams and tips are helpful and guide you to take your exam with success.

CCIE Routing and Switching v4.0 Configuration Practice Labs

Martin J. Duggan

Copyright© 2010 Pearson Education, Inc.

Published by:

Cisco Press

800 East 96th Street

Indianapolis, IN 46240 USA

All rights reserved. No part of this book may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or by any information storage and retrieval system, without written permission from the publisher, except for the inclusion of brief quotations in a review.

Printed in the United States of America

First Printing May 2010

ISBN-10: 1-58714-213-9

ISBN-13: 978-1-587-213-0

Warning and Disclaimer

This book is designed to provide information about the CCIE Routing and Switching version 4.0 lab exam. Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied.

The information is provided on an “as is” basis. The authors, Cisco Press, and Cisco Systems, Inc. shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this book or from the use of the discs or programs that may accompany it.

The opinions expressed in this book belong to the author and are not necessarily those of Cisco Systems, Inc.

Trademark Acknowledgments

All terms mentioned in this book that are known to be trademarks or service marks have been appropriately capitalized. Cisco Press or Cisco Systems, Inc. cannot attest to the accuracy of this information. Use of a term in this book should not be regarded as affecting the validity of any trademark or service mark.

Feedback Information

At Cisco Press, our goal is to create in-depth technical books of the highest quality and value. Each book is crafted with care and precision, undergoing rigorous development that involves the unique expertise of members from the professional technical community.

Readers’ feedback is a natural continuation of this process. If you have any comments regarding how we could improve the quality of this book, or otherwise alter it to better suit your needs, you can contact us through email at feedback@ciscopress.com. Please make sure to include the book title and ISBN in your message.

We greatly appreciate your assistance.

Corporate and Government Sales

Cisco Press offers excellent discounts on this book when ordered in quantity for bulk purchases or special sales. For more information, please contact:

U.S. Corporate and Government Sales 1-800-382-3419 corpsales@pearsontechgroup.com

For sales outside of the U.S. please contact: International Sales international@pearsoned.com



Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6367)
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems, Inc.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Europe Headquarters
Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: +31 0 800 020 0791
Fax: +31 0 20 357 1100

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

©2007 Cisco Systems, Inc. All rights reserved. CCVP, the Cisco logo, and the Cisco Square Bridge logo are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn is a service mark of Cisco Systems, Inc.; and Access Registrar, Aironet, BPX, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Enterprise/Solver, EtherChannel, EtherSwitch, Fast Step, Follow Me Browsing, FormShare, GigaDrive, GigaStack, HomeLink, Internet Quotient, IOS, IP TV, IQ Expertise, the IQ logo, IQ Net Readiness Scorecard, iQuick Study, LightStream, Linksys, MeetingPlace, MGX, Networking Academy, Network Registrar, Packet, PIX, ProConnect, RateMUX, ScriptShare, SlideCast, SMARTnet, StackWise, The Fastest Way to Increase Your Internet Quotient, and TransPath are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0809R)