CHAPTER 1

Practice Lab 1

Each lab has a time constraint of eight hours and a point scale weighting of 100; you will need to score at least 80 marks to pass. The lab has been designed to challenge you in areas that you will find in the real exam with each lab having a distinct theme to enhance your study plan; Routing Information Protocol (RIP V2 is the theme of Lab 1).

You will, of course, find the old favorites such as BGP, DLSw+, and Voice but a complete understanding of RIP V2 will earn you extra points in this lab.

Aim to adhere to the time limit on this lab on the initial run through and then either score yourself at this point or continue until you feel you have met all the objectives. Keep a note of your score to plot your progress throughout the book and remember you are aiming to improve your technical knowledge, speed, and examination technique.

If you find that you complete all the configuration tasks within the time limit, congratulations, you are a quick on the keyboard but will you achieve the desired results? If time allows, get into the habit of going back through the questions and ensuring that you have answered them down to the letter. If you are unsure, turn to the Lab 1 "Ask the Proctor" section but try not to use this too often as you will find that real-life proctors do not like to give anything away. However, throughout this book, it can be used as a handy tool to provide assistance and clues to ensure you are working on the correct solution for the question. Unfortunately you won't have this luxury in your real exam.

You might find the questions misleading or vague but if you re-read the information given and analyze the scenario, you will find that you have been given sufficient information to successfully solve the problem.

To assist you, initial and final solutions are provided for the entire lab including configurations and common show command outputs from all the devices in the topology on the accompanying CD. The aforementioned "Ask the Proctor" section is included at the end of the lab, which gives you clues, if required, followed by the lab debrief that analyzes each question showing you what was required and how to achieve the desired results. Finally, you will find handy references should you require additional study information.

You will now be guided through the equipment requirements and pre-lab tasks in preparation for taking Practice Lab 1.

Equipment List

You need the following hardware and software components to begin Lab 1.

- Eight routers are required loaded with Cisco IOS Software Release 12.2-16
- Enterprise image and the minimum interface configuration as documented in Table 1-1:

Table 1-1 *Interfaces Required per Router*

Router	Ethernet Interface	Serial Interface	BRI Interface	Voice	ATM Interface
R1	1	2	1	1 X FXS	-
R2	1	1	-	-	-
R3	1	-	-	-	-
R4	1	1	1	1 X FXS	-
R5	1	3	-	-	1
R6	2	1	-	-	1
R7	1	-	-	-	-
R8	1	-	-	-	-

NOTE Lab 1 was produced with Routers R1, R2, R3, R4, R7, and R8 using 2600s and R5 and R6 using 7200s.

- One Switch 3550 with Cisco IOS Software Release 12.1(12c) enterprise: c3550-i5q3l2-mz.121-12c.EA1.bin

Setting Up the Lab

Feel free to use any combination of routers as long as you fulfill the topology diagram as shown in Figure 1-1. It is not compulsory to use the same model of routers, but this will make life easier should you want to load configurations directly from the CD-ROM into your own devices.

NOTE For each lab in the book you will have a set of initial configuration files that can be different from each other. Notice that some interfaces will not have the P address preconfigured, because you will either not be using that interface on that specific lab or because you will need to work on this interface through the exercise. The initial configurations can be found on the CD-ROM and should be used to preconfigure your routers and switch before the lab starts.

If you use the same equipment as used to produce the lab, you can simply paste the configurations into your own equipment; if not, just configure your own equipment accordingly using the information supplied within the initial configurations.

Labs 1 through 3 in this book have been completed using 100-Mbps Fast Ethernet interfaces so if you have a mix of 10- and 100-Mbps Ethernet interfaces, adjust the bandwidth statements on the relevant interfaces to keep all interface speeds common. This will ensure that you do not get unwanted behavior because of differing IGP metrics.

Lab Topology

Practice Lab 1 uses the topology as outlined in Figure 1-1, which you need to create using the switch, Frame Relay, ATM, and ISDN information that follows.

Figure 1-1 *Lab 1 Topology Diagram*

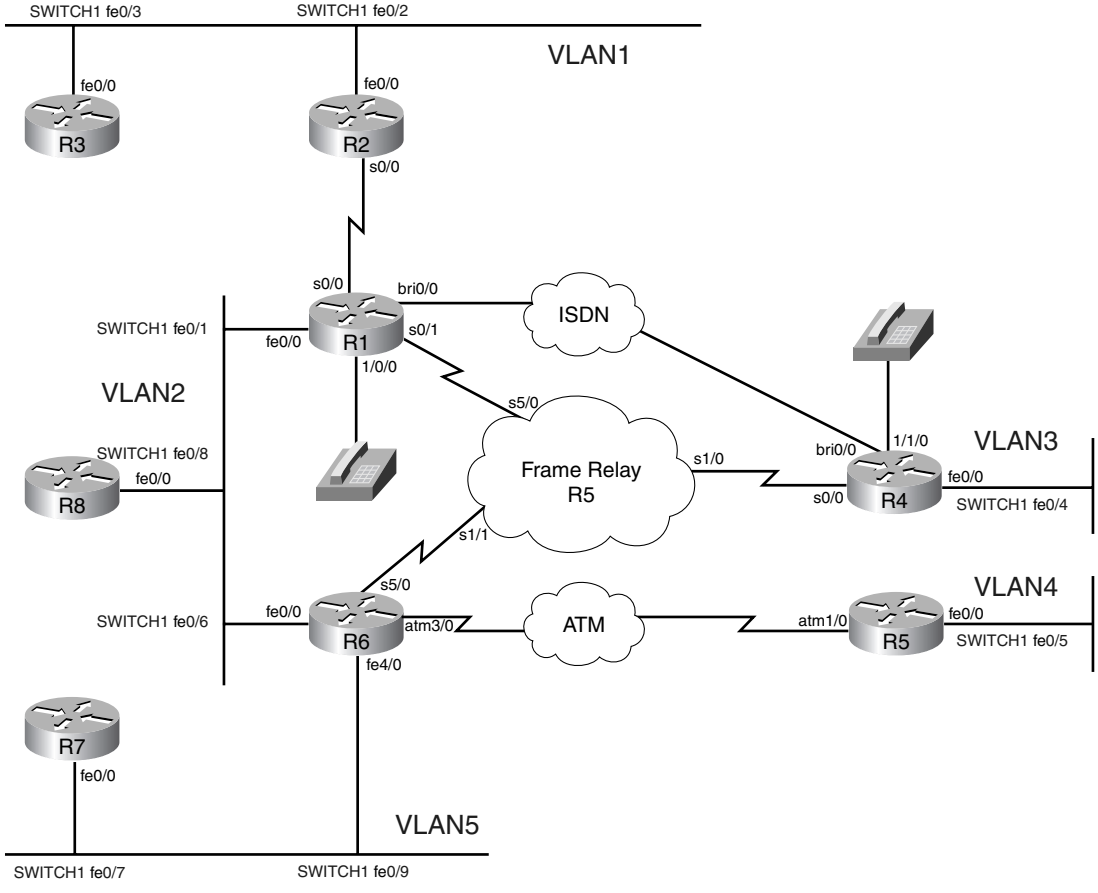

Cabling Instructions

Follow the cabling requirements as outlined in Figure 1-2 and Table 1-2 to connect your routers to the switch.

Figure 1-2 *3550 Cabling Diagram*

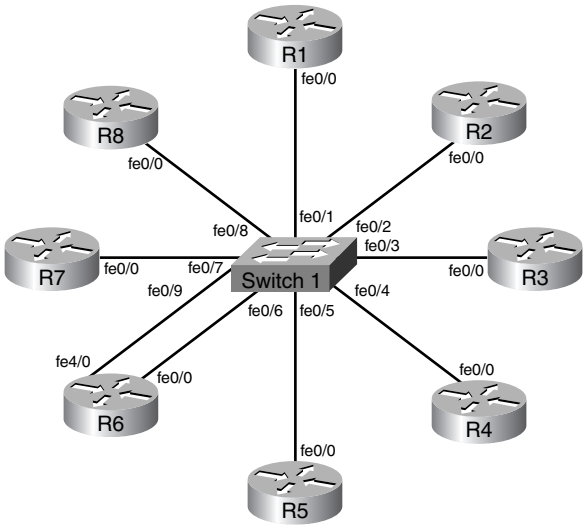

Table 1-2 *3550 Cabling Guide*

Ethernet Cabling	Switch1 Port Number
R1-Fast Ethernet0/0	Port 0/1
R2-Fast Ethernet0/0	Port 0/2
R3-Fast Ethernet0/0	Port 0/3
R4-Fast Ethernet0/0	Port 0/4
R5-Fast Ethernet0/0	Port 0/5
R6-Fast Ethernet0/0	Port 0/6
R7-Fast Ethernet0/0	Port 0/7
R8-Fast Ethernet0/0	Port 0/8
R6-Fast Ethernet4/0	Port 0/9

Frame Relay Switch Instructions

The Frame Relay portion of the lab is achieved by following the physical connectivity using R5 as a Frame Relay switch as shown in Figure 1-3.

Figure 1-3 *Frame Relay Switch Physical Connectivity*

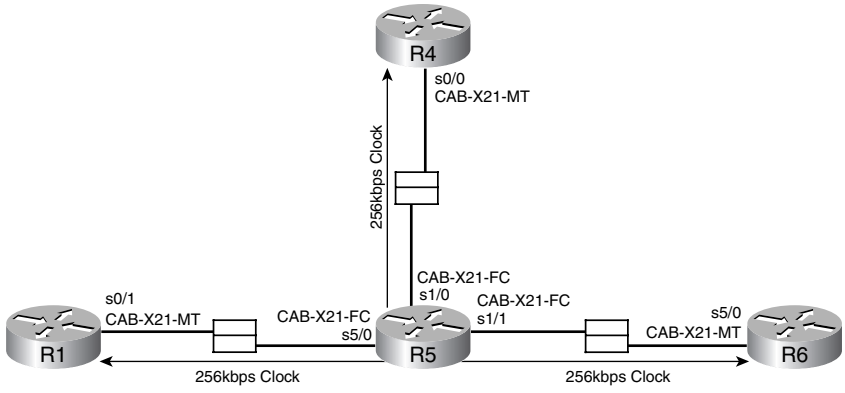

The physical Frame Relay connectivity (after configuration) will represent the logical Frame Relay network as shown in Figure 1-4.

Figure 1-4 *Frame Relay Switch Logical Connectivity*

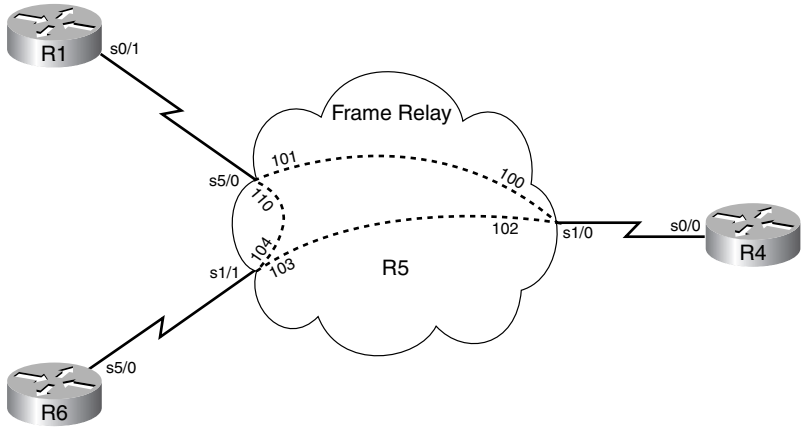

Configure one of your routers as a Frame Relay switch or have a dedicated router purely for this task. The first three lab scenarios use R5 to form the Frame Relay switch and a fully meshed environment is configured between R1-R4-R6, so pay attention in the lab to which PVCs are actually required. Keep the encapsulation and Local Management Interface (LMI) settings to default for this exercise, but experiment with the settings outside the labs.

Keep your DCE cables at the Frame Relay switch end for simplicity and provide a clock rate of 256 kbps to all links. Should you require detailed information on how to configure one of your routers as a Frame Relay switch, this information can be found in Appendix A, "Frame Relay Switch Configuration."

NOTE The Frame Relay switch configuration for R5 is supplied on the CD-ROM, if required.

ATM Switch Instructions

The ATM portion of the lab is achieved by following the physical connectivity between R5 and R6 as shown in Figure 1-5.

Figure 1-5 *ATM Physical Connectivity*

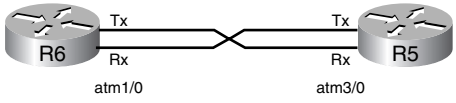

The physical ATM connectivity will, after configuration, represent the logical ATM network as shown in Figure 1-6.

Figure 1-6 *ATM Logical Connectivity*

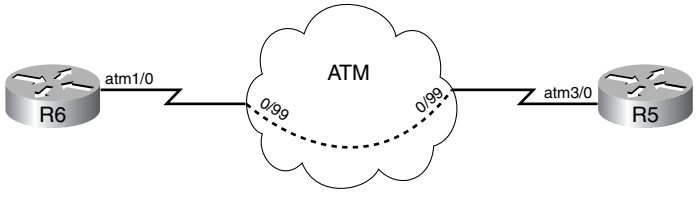

The ATM connectivity in Labs 1-5 will be provided by back-to-back connections between R6 and R5 over E3 ATM interfaces (you could also use a LightStream or whichever back-to-back flavor of ATM you have available). Configure the PVCs as requested during the Lab exercise. If you are using a LightStream to provide your ATM connectivity and require information on how to set this up, this information can be found in Appendix B, "LS1010 ATM Switch Configuration."

Serial Back-to-Back Instructions

R1 and R2 are connected back-to-back with serial cables as shown in Figure 1-7. Ensure that the DCE cable is connected to R1 and generate a 2 Mbps clock from this point if using X21 cables as shown or reduce this to suit your own serial interfaces such as 1.5 Mbps for T1 connectivity.

Figure 1-7 *Serial Connectivity*

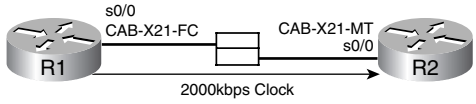

ISDN Instructions

Connect R1 and R4 into either ISDN lines or an ISDN simulator. It is preferable that the ISDN supports CLI. Reconfigure the numbers as required if you are using live ISDN lines.

The lab has been produced using BRI S/T interfaces on R1 and R4 as shown in Figure 1-8.

Figure 1-8 *ISDN Connectivity*

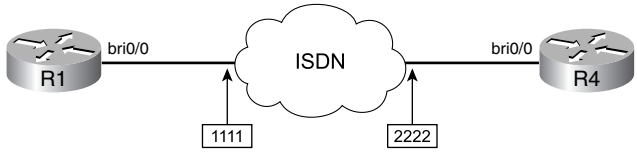

IP Address Instructions

Configure the IP addresses as shown in Figure 1-9 or load the initial router configurations for Lab 1 that can be found on the CD-ROM. If manually configuring, ensure you include the following loopback addresses:

R1 lo0 10.1.1.1/28
R4 lo0 10.4.4.4/29
R5 lo0 10.5.5.5/30
R6 lo0 10.6.6.6/29
R7 lo0 10.7.7.7/28
R8 lo0 10.8.8.8/32

Figure 1-9 *IP Addressing Diagram*

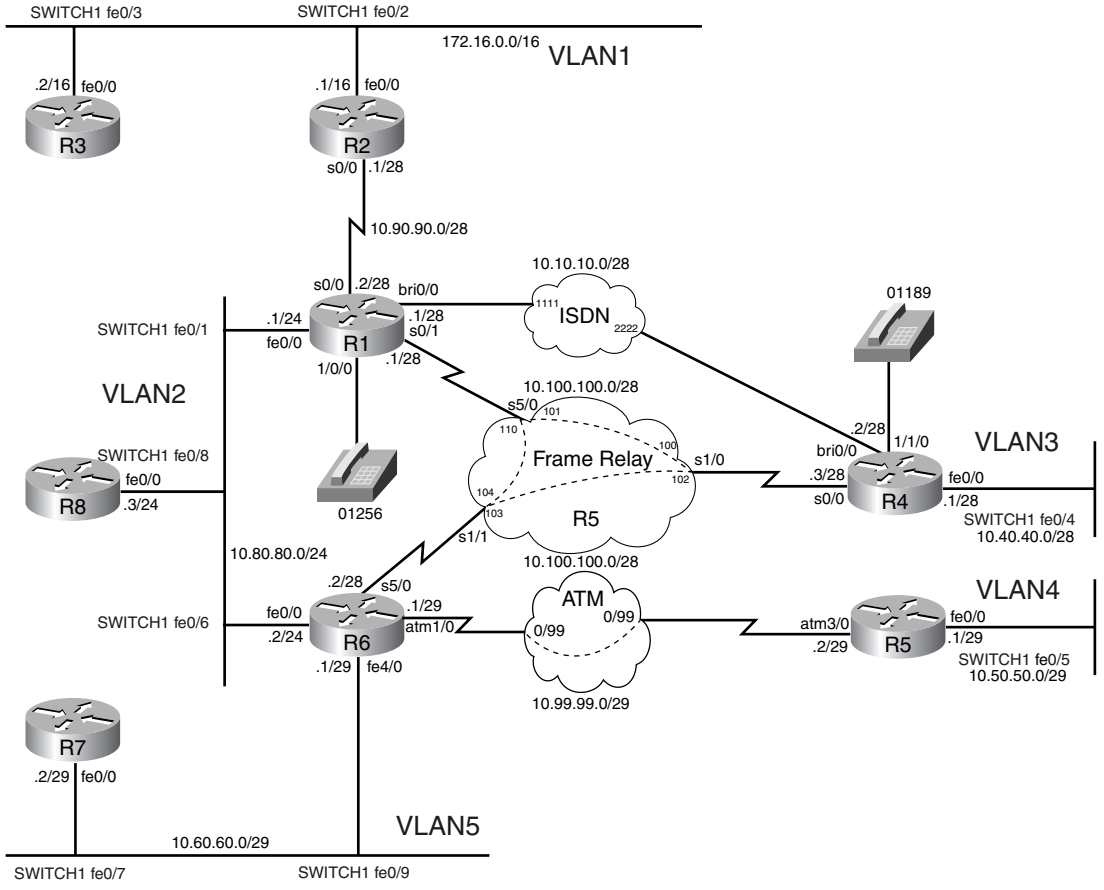

Pre-Lab Tasks

- Build the lab topology as per Figure 1-1 and Figure 1-2.
- Configure your chosen Frame Relay switch router to provide the necessary data-link control identifiers (DLCIs) as per Figure 1-4 or load the Frame Relay switch configuration from the CD-ROM.

- Configure the IP addresses on each router as shown in Figure 1-9 and add the loopback addresses (do not configure the Frame Relay or ATM IP addresses yet as you will need to select interface types within the lab beforehand); alternatively, you can load the initial configuration files from the CD-ROM.
- Configure passwords on all devices for console and vty access to "cisco" if not loading the initial configuration files.
- If you find yourself running out of time, choose questions that you are confident you can answer correctly. Another approach would be to 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 eight hours.

General Guidelines

- Please read the whole lab before you start.
- Do not configure any static/default routes unless otherwise specified.
- Use only the DLCIs and ATM PVCs provided in the appropriate figures.
- Ensure full IP visibility between routers for ping testing/telnet access to your devices.
- 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/univercd/home/home.htm

NOTE Consider accessing only the preceding URL, not the entire Cisco.com website. If you will be allowed to use online documentation during your CCIE lab exam, it will be restricted.

Practice Lab 1

You will now be answering questions in relation to the network topology as shown in Figure 1-10.

Figure 1-10 *Lab 1 Topology Diagram*

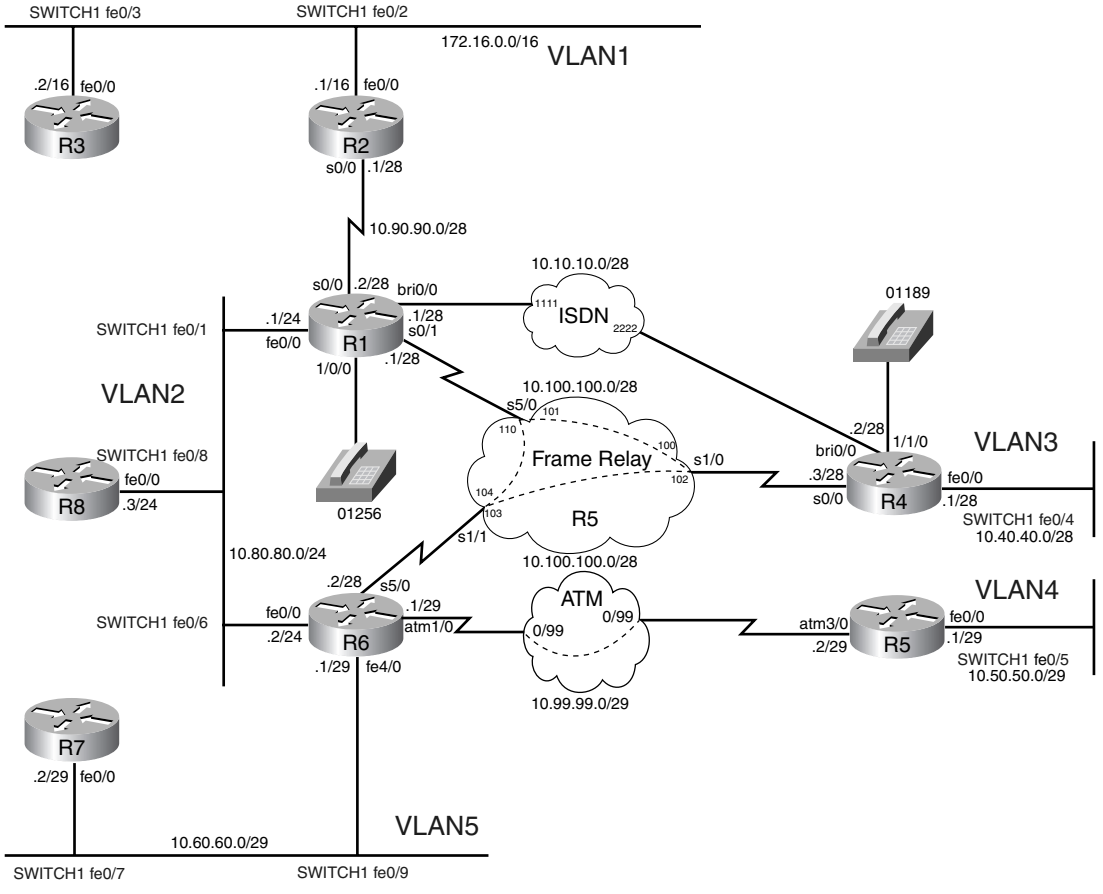

Section 1: Bridging and Switching (15 Points)

Section 1.1: Frame Relay Configuration (6 points)

- Configure the Frame Relay portion of the network as shown in Figure 1-11; ensure that DLCIs 110 and 104 between R1-R6 are not used.
- The routers are to be on the same subnet and should be configured with subinterfaces.

Figure 1-11 *Frame Relay Diagram*

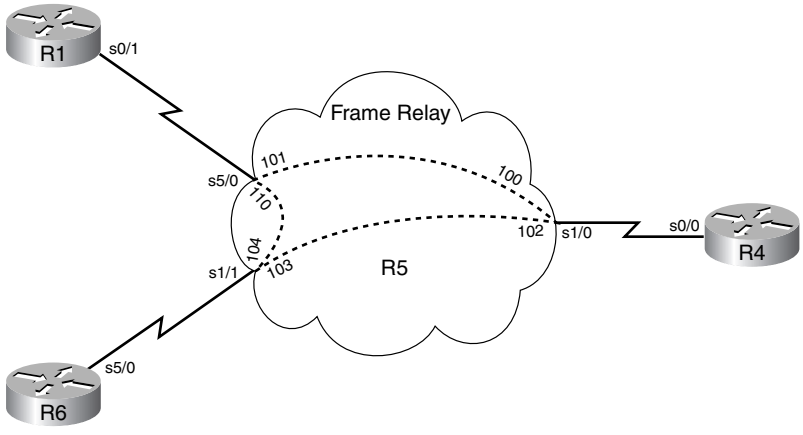

Section 1.2: 3550 LAN Switch Configuration (6 Points)

- Configure VLAN numbers, VLAN names, and port assignment as per the topology diagram as shown in Figure 1-10.
- There is to be a host connected on interface 0/16 in the future; the network administrator requires that this host is authenticated by a radius server before access to the switch is granted. The radius server is to be located on the IP address 172.16.100.100 with the key **radius14.**
- Ensure the switch is reachable via Telnet to the IP address of 10.80.80.8/24.

Section 1.3: ATM Configuration (3 Points)

- Configure the ATM network as shown in Figure 1-12.
- Use a subinterface on R6 for the ATM matching the VCI number and ensure the latest method of PVC configuration is used on this router. For R5 ATM, use the physical interface and legacy PVC configuration; after you have configured your Layer 2 information, you may then add the Layer 3 addresses.
- Do not rely on inverse Address Resolution Protocol (ARP).

Figure 1-12 *ATM Diagram*

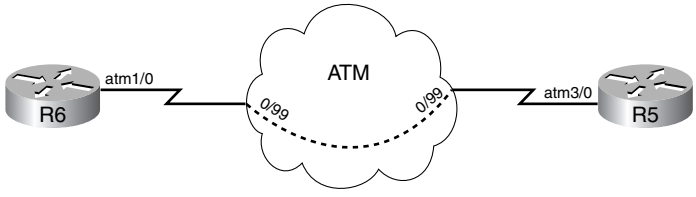

Section 2: IP IGP Protocols (28 Points)

Configure the IP routing as in Figure 1-13 and redistribute protocols to ensure full IP visibility between routers. Advertise all router networks within the appropriate routing protocol.

Section 2.1: RIP (16 Points)

- On all RIP routers, ensure that **version 2** is used under the process.
- Ensure that VLSM is supported on advertisements between all RIP routers.
- Add a loopback interface with the address of 60.60.60.1/24 onto R3 and advertise this out to R2 but ensure that it is not seen by the rest of your network; do not perform any configuration on R2 or R1.
- Configure R3 to unicast its RIP routing updates to R2. Do not use the **neighbor** command to achieve this but consider using other IP features to aid you.
- Ensure that VLAN2 is advertised to the RIP domain as a /28 network. Do not use either RIP or EIGRP features to accomplish this. You can, however, configure R6.

Section 2.2: EIGRP (5 Points)

- R8 is very low on memory and CPU resource; accommodate this information within the configuration on R8.
- Configure R8 to have an EIGRP hello interval of 25 seconds on its FastEthernet0/0 interface.

Figure 1-13 *IP IGP Diagram*

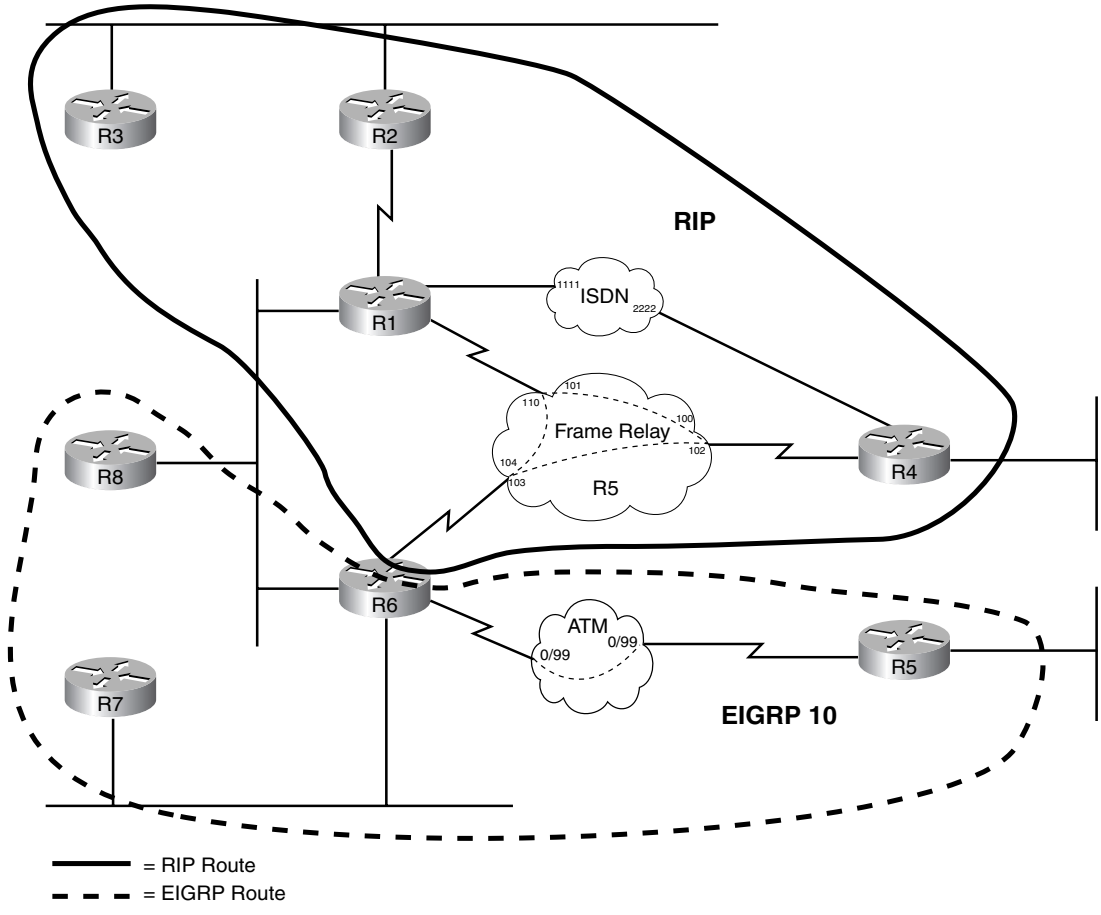

Section 2.3: Redistribution (7 Points)

- Redistribute IGP protocols to ensure full IP visibility between all routers.
- As a safety precaution, ensure that R6 can not learn the EIGRP routes it previously advertised into the RIP domain back from R4.

Section 3: ISDN (8 Points)

- Ensure that VLAN3 and R4 Lo0 are accessible from R1 and beyond should the Frame Relay network fail either physically or logically. If VLAN3 and R4 Lo0 networks are restored while the ISDN line is active, ensure that traffic is routed over the Frame Relay network to these destinations immediately.
- Configure R1 so that if half of the ISDN traffic to R4 is of an unacceptable quality, the line is automatically disconnected.
- Allow only R1 to dial into R4. Do not use any PPP feature in your solution.
- Do not allow the ISDN to flap if the Frame Relay network goes up and down; only allow the line to be dropped if the Frame Relay connectivity is deemed to be reliable for 90 seconds.

Section 4: EGP Protocols (17 Points)

- Configure BGP, as shown in Figure 1-14, with the following peering: R3–R2, R8–R2, R6–R2, R6–R8, R7–R8, R6–R5. Ensure the most suitable interfaces are used to maintain resilience for BGP peering (except for R2 and R3, use 172.16.0.0/16 addresses for all peering to and from these routers).
- Ensure minimal configuration on R8.
- Inject the following networks into BGP via new loopback interfaces:

  R3: 20.200.200.1/24 and 20.20.20.1/24

  R5: 20.20.20.1/24 and 200.20.20.1/24

  R7: 30.30.30.30/29
- Ensure that R6 BGP routing table prefers to use AS1000 for network 20.20.20.0/24; do not use BGP weight, BGP local preference, MED, neighbor metric related statements, metric manipulation, summarization, or pre-pending to achieve this. Perform configuration on R6 only.
- All BGP speakers are to be able to communicate with all advertised BGP networks.

Figure 1-14 *IP EGP Diagram*

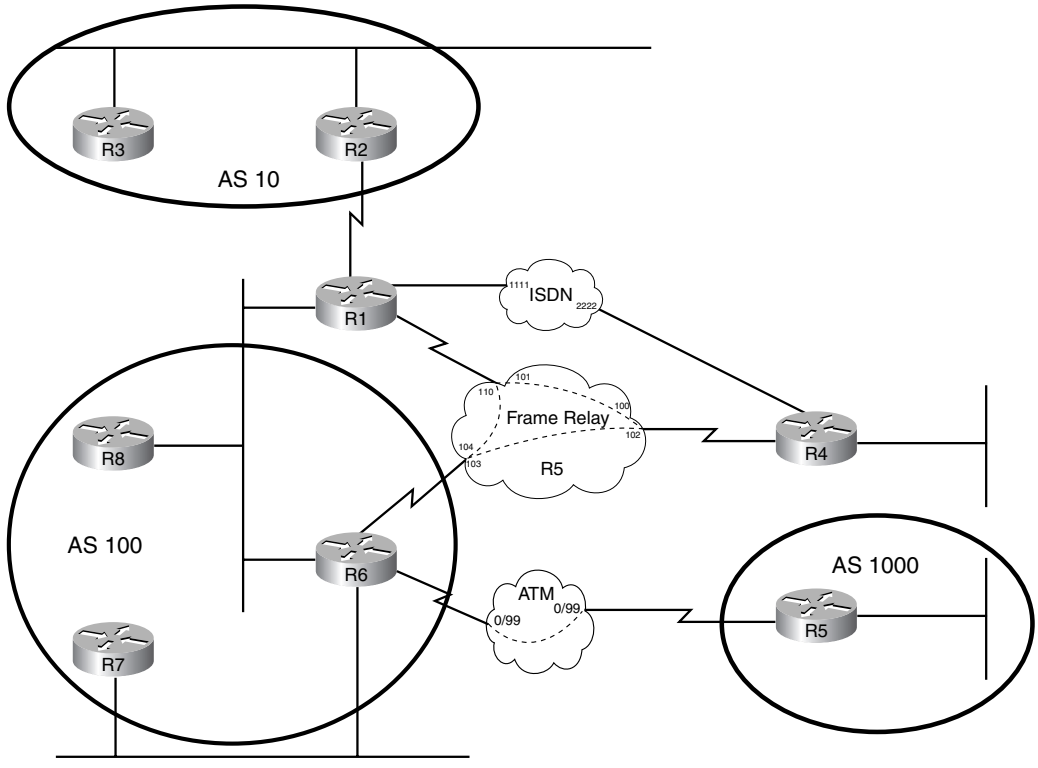

Section 5: Voice (6 Points)

- Both phones should be able to ring each other using the numbers supplied. Use the most efficient method of transporting the voice from R1 to R4.
- Ensure that voice is still available if the main connection between R1 and R4 fails.
- Make the phone on R1 also answer calls if 01962 is dialed on the R4 connected handsets; do not use number expansion to achieve this.

Section 6: DLSw+ (4 Points)

- Configure DLSw+ between VLAN2 and VLAN4; use routers R8 and R5. Peer from Lo0 on R8 and the VLAN4 interface on R5; ensure that R8 can accept DLSw+ connections from only unknown TCP peers.
- Set up a one-line filter, which only allows common SNA traffic to egress from R5 VLAN4 into the DLSw+ network.

Section 7: IOS and IP Features (10 Points)

- R2 is sited in a shared data center; make the serial link back into R1 as secure as possible at Layer 2.
- Ensure traffic from VLAN4, including any attached router interfaces to VLAN4, is hidden behind R5 lo0 address when directed toward all external router networks.
- A router is to be installed onto VLAN4 in the future. This router will have a default configuration, so allow R6 to assist dynamically to aid the configuration process. The router will require an IP address of 10.50.50.6 and should load a configuration file called R9-config from a fictitious TFTP server on 172.16.0.59.

Section 8: QoS (8 Points)

- Achieve maximum quality of voice calls by ensuring the real-time packet interval of 10 ms is not exceeded. Do not use RSVP in your solution.
- To reduce the packet fragmentation in your network, allow R5 to determine appropriate fragmentation requirements when TCP sessions are originated from it to any part of the network.

Section 9: Multicast (4 Points)

- Enable your network to allow hosts on VLAN4 to receive and send multicast traffic from and to VLAN2; only perform configuration on R5 and R6 using PIM sparse dense mode.
- Configure R6 to respond to pings from R5 to the multicast address of 224.4.4.4.
- Do not allow R5 to fully participate in the PIM process by not allowing it to become a neighbor, but do allow any IGMP messages generated by hosts on VLAN4 to be received by R6.

Practice Lab 1: "Ask the Proctor"

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

Section 1.1: Frame Relay Configuration

Q: If I don't configure the PVC between R1-R6, surely it won't be used?
A: You have to make sure that the PVC is not used.

Q: But if it's not configured, surely it won't be used?
A: Use the command **show frame relay pvc**; if the PVC shows any input packets or output packets, you have not answered the question. You can also use the command **show frame map.**

Q: It says the routers must use subinterfaces. Can I use point-to-point everywhere?
A: This does not address the question.

Q: Can I use two separate point-to-point subinterfaces from R4 out to individual point-to-point subinterfaces on R1 and R6?
A: If the routers were not on a common subnet you could, but this is not the case.

Q: Can I put a point to multi-point interface on R4 and point-to-point interfaces on R1 and R6?
A: Yes.

Q: Do you want me to configure "**broadcast**" under my Frame Relay interfaces?
A: You need to determine what type of traffic/protocols will use the Frame Relay network.

Section 1.2: 3550 LAN Switch Configuration

Q: Do I have to configure the switch as a VTP server?
A: The questions do not ask you to do this and there is only one switch in the network.

Q: Why can't I rename VLAN1?
A: You can't so don't waste time on this.

Q: Can I leave my switch ports in auto mode?
A: The questions do not specifically ask you to configure speed and duplex, but you should do this to avoid any mismatches that could cause you problems.

Q: For the future port 0/16, can I leave it in VLAN1?
A: There is no specific VLAN stated in the question, so yes.

Section 1.3: ATM Configuration

Q: What kind of subinterface do you want on R6?
A: There are only two routers connected over ATM; choose a suitable subinterface type to reflect this.

Q: When you say latest method of configuration, do you mean using the **pvc** command?
A: Yes.

Q: When you say legacy configuration, do you mean using the **map-list** command?
A: Yes.

Q: What ATM encapsulation should I use?
A: This is a back-to-back configuration so use a suitable type of your choice.

Q: Do you want me to configure "**broadcast**" under my ATM interfaces?
A: You need to determine what type of traffic/protocols will use the ATM network.

Section 2.1: RIP

Q: Can I use a **distribute-list** to stop the advertisement of network 60.60.60.0/24?
A: You cannot configure R2 or R1; R2 must see the network so you cannot configure a distribute-list on R3 either.

Q: Surely the only way to make RIP unicast routing updates is to use the **neighbor** command and **passive-interface**?
A: There is another way of forcing this; try to use other IP features if you find that you can not accomplish this using RIP features.

Q: Is it acceptable to use a NAT list to convert my RIP multicast into a unicast?
A: If you answer the question effectively, this is acceptable.

Q: I used NAT and now my connectivity is lost between R2 and R3. Is this supposed to happen?
A: You need to maintain full IP connectivity between all routers as well as answering specific questions.

Q: Can I use summarization to advertise VLAN2 into the /28 RIP V2 domain?
A: The question clearly states no RIP or EIGRP features may be used.

Q: Can I run ospf on my router and redistribute this into RIP with an ospf summary of VLAN2?
A: The question states that no RIP feature may be used; redistribution is a feature of RIP.

Q: Am I permitted to add a secondary address to my configuration?
A: If this achieves the desired result, yes.

Q: I am having problems connecting to VLAN2 from R1 after answering the VLAN2 question. Is this expected?
A: If you experience any connectivity issues, you should investigate and rectify them.

Section 2.2: EIGRP

Q: Can I just use passive interfaces to reduce the CPU processing?
A: Passive interfaces with no requirements for adjacencies will reduce CPU processes, but a superior way to enable this within EIGRP exists.

Q: If I use the **neighbor** command, won't that be sufficient as I will no longer be multicasting my EIGRP "hellos"?
A: This will not be sufficient.

Q: I have modified my hello timer on R8 and now my routes are flapping. Is this acceptable?
A: No, you need to remember that R8 will have formed a neighbor relationship on VLAN2; you should ensure that all routers are configured in the same manner.

Q: I have configured R6 and R8 with identical hello intervals but I find that my routes are still flapping. Is this acceptable?
A: No, you should maintain a reliable neighbor relationship with R4; consider adjusting other EIGRP parameters.

Section 2.3: Redistribution

Q: I believe I can complete my redistribution without any filtering. Is this acceptable?
A: To avoid suboptimal routing and potential routing loops, it is always good practice to filter when redistributing.

Q: But no way exists that I could create suboptimal routes in this scenario?
A: Correct. Use this question as practice.

Section 3: ISDN

Q: Can I use **backup interface** on R1 to protect the Frame Relay network?
A: The question specifically asks you for visibility of VLAN3 and R4 lo0 networks, Further networks would be available if you were to use the backup interface. This command would also only work if there was a physical Layer 1 problem with the Frame Relay connection.

Q: Can I use floating static routes for VLAN3 and R4 lo0 networks?
A: You are not permitted to use any static routes in any part of the lab.

Q: I have managed to get the ISDN to dial out if I loose visibility of VLAN3 and Lo0 but the line will not stay down. Is this acceptable?
A: If the Frame Relay network is restored, your ISDN line should eventually be disconnected.

Q: Do you want CHAP or PAP configured over the ISDN so only R1 dials into R4?
A: The question does not specifically ask for CHAP or PAP or give sufficient information to configure either.

Q: Can I use my **dialer idle-timeout** set to 90 seconds to ensure the ISDN does not flap?
A: Adjusting the **dialer idle-timeout** is not sufficient.

Section 4: EGP Protocols

Q: Is it sufficient if I peer from connected interfaces as default?
A: Where possible, you should peer using the most resilient method as instructed.

Q: I have just configured my routers to peer from their loopback interfaces and now not all of my BGP neighbors are showing up. Is this acceptable?
A: No, you need to remember the rules for EBGP peering.

Q: To minimize the configuration on R8, can I leave out statements such as **update-source**?
A: No, find a way to maintain the required features but cut down on the size of your configuration.

Q: Can I advertise my new loopback interfaces as I see fit.
A: Use the most appropriate method of advertising your new networks.

Q: Can I change router IDs to manipulate the path selection on R6?
A: You can try anything not listed in the question if it achieves the desired results.

Q: If I have the BGP routes in my BGP tables, surely this is sufficient to prove my BGP is functioning correctly and I have full IP visibility?
A: Potentially yes, however, it would be prudent to perform extended pings or by using **traceroute** to be 100-percent certain.

Section 5: Voice

Q: Because voice should be available over the Frame Relay and ISDN, is it acceptable to run just Voice over IP?
A: A more efficient means of transportation of Voice over the Frame Relay network exists.

Q: If I can't use **num-exp**, surely I can't get the other phone to ring?
A: A method is available.

Section 6: DLSw+

Q: Can I set up multiple **remote-peer** statements on R8 from other loopback interfaces within the network for future peering?
A: You need to allow for any unknown future peering.

Q: With my filter, I can't allow for every form of SNA SAP?
A: Allow for the most common SAPs but keep your list to one line.

Section 7: IOS and IP Features

Q: Do you require IPsec configured between R1 and R2?
A: The question refers to Layer 2; IPsec is a Layer 3 protocol.

Q: Can I use access lists to just allow R2 to communicate with R1?
A: This is again Layer 3.

Q: Do you want me to authenticate my routing updates to ensure security?
A: You need a Layer 2 solution.

Q: Do you want me to use a single NAT instance to hide behind my loopback interface?
A: You may find that you require multiple instances to answer the question effectively.

Q: Surely the new router will have at least been configured to pick up a DHCP address?
A: The new router is out of the box with a factory configuration only.

Q: Do you actually want the configuration of R9 to be stored on R6?
A: No, the configuration is to be held on a fictitious TFTP server.

Section 8: QoS

Q: Do you just want me to configure custom queuing for voice?
A: No, you must ensure that voice is transmitted within 10 ms intervals so the quality is not impaired.

Q: Do you want the QoS for the voice when it is transmitted over Frame Relay or ISDN?
A: It is your decision how and where you activate QoS, sufficient information is available in the paper for you to make an informed decision.

Q: For R5 packet fragmentation, is it acceptable to work out the smallest MTU in the network and configure this on all interfaces on R5?
A: Although this might reduce fragmentation at other points within the network, it does not allow R5 to determine the fragmentation itself.

Section 9: Multicast

Q: Do you want me to configure any multicast parameters on my switch?
A: The question states only R5 and R6 should be configured.

Q: Can I create a standard access-list blocking PIM from R5 and R6?
A: A more elegant PIM filtering method of achieving this exists.

Practice Lab 1 Debrief

The lab debrief section will now analyze 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 your test.

Section 1: Bridging and Switching (15 Points)

Section 1.1: Frame Relay Configuration (6 points)

- *Configure the Frame Relay portion of the network as shown in Figure 1-8; ensure that DLCIs 110 and 104 between R1-R6 are not used.*

The question clearly states that DLCIs 110 and 104 are not to be used; you must, therefore, disable **inverse-arp** on the routers. It is good practice to ensure that all routers do not rely on **inverse-arp** so if you have configured **no frame-relay inverse-arp** under routers R1,R4 and R6 serial interfaces 0/0, you have scored 2 points.

If you experience difficulties and can not clear any dynamic map entries, reload your routers to remove these, a drastic measure but every point counts.

- *The routers are to be on the same subnet and should be configured with subinterfaces.*

R4 will need to be a multipoint subinterface to accommodate both R1 and R6 on the same subnet; R1 and R6 only have PVCs to R4, hence, they will require point-to-point subinterfaces. R4 will require manual **frame-relay map** statements pointing to both R1 and R6 as inverse arp is disabled. The maps require the **broadcast** keyword as RIP will multicast the routing updates over the PVCs. It should be apparent that when RIP is run over a multipoint interface, split horizon will be enabled by default and routing updates from R6 into R1 will never be propagated by the hub router R4 because of the rule of not advertising a network that was received on the same interface; R4 will, therefore, require **no ip split-horizon** configured under its Frame Relay interface. If you have configured all items correctly as in Example 1-1 through Example 1-3, you have scored 4 points, unfortunately no marks if you have omitted anything.

NOTE For clarity only, the required configuration details will be listed to answer the specific questions instead of full final configurations.

Example 1-1 *R4 Initial Frame Relay Solution Configuration*

```
interface Serial0/0
 no ip address
 encapsulation frame-relay
 no frame-relay inverse-arp
!
interface Serial0/0.1 multipoint
 ip address 10.100.100.3 255.255.255.240
 no ip split-horizon
 frame-relay map ip 10.100.100.1 100 broadcast
 frame-relay map ip 10.100.100.2 102 broadcast
```

Example 1-2 *R1 Initial Frame Relay Solution Configuration*

```
interface Serial0/1
 no ip address
 encapsulation frame-relay
 no frame-relay inverse-arp
!
interface Serial0/1.101 point-to-point
 ip address 10.100.100.1 255.255.255.240 frame-relay interface-dlci 101
```

Example 1-3 *R6 Initial Frame Relay Solution Configuration*

```
interface Serial5/0
 no ip address
 encapsulation frame-relay
 no frame-relay inverse-arp
!
interface Serial5/0.103 point-to-point
 ip address 10.100.100.2 255.255.255.240
 frame-relay interface-dlci 103
```

Section 1.2: 3550 LAN Switch Configuration (6 Points)

- *Configure VLAN numbers, VLAN names, and port assignment as per the topology diagram as shown in Figure 1-10.*

The switch in this instance is isolated but you can still use the default mode of VTP Server. From the VLAN database, add the required VLANs and name them accordingly; you

should note that you can not change the VLAN name of VLAN1. You must ensure that the port speed and duplex is fixed to 100 Mbps and full duplex, if your routers support this; leaving your ports in auto mode could cause connectivity problems. If you have configured these items correctly as in Example 1-4, you have scored 2 points.

Example 1-4 *3550 Switch1 Initial Configuration*

```
Switch1#vlan database
Switch1(vlan)#vlan 2 name VLAN2
VLAN 2 modified:
    Name: VLAN2
Switch1(vlan)#vlan 3 name VLAN3
VLAN 3 modified:
    Name: VLAN3
Switch1(vlan)#vlan 4 name VLAN4
VLAN 4 modified:
    Name: VLAN4
Switch1(vlan)#vlan 5 name VLAN5
VLAN 5 modified:
    Name: VLAN5
Switch1(vlan)#exit
APPLY completed.
Exiting....
interface FastEthernet0/1
 switchport access vlan 2
 switchport mode access
 no ip address
 duplex full
 speed 100
!
interface FastEthernet0/2
 switchport mode access
 no ip address
 duplex full
 speed 100
!
interface FastEthernet0/3
 switchport mode access
 no ip address
 duplex full
 speed 100
!
interface FastEthernet0/4
 switchport access vlan 3
 switchport mode access
 no ip address
 duplex full
 speed 100
!
interface FastEthernet0/5
 switchport access vlan 4
 switchport mode access
 no ip address
```

continues

Example 1-4 *3550 Switch1 Initial Configuration (Continued)*

```
 duplex full
 speed 100
!
interface FastEthernet0/6
 switchport access vlan 2
 switchport mode access
 no ip address
 duplex full
 speed 100
!
interface FastEthernet0/7
 switchport access vlan 5
 switchport mode access
 no ip address
 duplex full
 speed 100
!
interface FastEthernet0/8
 switchport access vlan 2
 switchport mode access
 no ip address
 duplex full
 speed 100
!
interface FastEthernet0/9
 switchport access vlan 5
 switchport mode access
 no ip address
 duplex full
 speed 100
```

NOTE The VLAN configuration is completed under **vlan database.**

- *There is to be a host connected on interface 0/16 in the future; the network administrator requires that this host is authenticated by a radius server before access to the switch is granted. The radius server is to be located on the IP address 172.16.100.100 with the key* **radius14.**

This question calls for 802.1X Authentication before a port is granted access to the switch and network. If configured correctly as in Example 1-5, you have scored 3 points.

Example 1-5 *802.1X Switch Configuration*

```
aaa new-model
aaa authentication dot1x default group radius
!
interface FastEthernet0/16
```

Example 1-5 *802.1X Switch Configuration (Continued)*

```
 switchport mode access
 no ip address
 dot1x port-control auto
!
radius-server host 172.16.100.100 auth-port 1812 key radius14
```

- *Ensure the switch is reachable via Telnet to the IP address of 10.80.80.8/24.*

Configure VLAN2 with the IP address of 10.80.80.8 255.255.255.0. The switch will also need a **default-gateway** configured; you could use 10.80.80.2 or 10.80.80.1 here. The previous question requires that you enable AAA. Enabling AAA prompts you for a username when you telnet to the switch from one of your routers. To ensure typical access to the preconfigured line and to ensure that the enable password is used for telnet access to the switch, you should add the **aaa authentication login default enable** authentication configuration onto the switch.

Example 1-6 *Switch1 Management IP Configuration*

```
aaa authentication login default enable
enable password cisco
!
interface Vlan2
 ip address 10.80.80.8 255.255.255.0
!
ip default-gateway 10.80.80.2
!
line con 0
 password cisco
line vty 0 15
 password cisco
```

Section 1.3: ATM Configuration (3 Points)

- *Configure the ATM network as shown in Figure 1-12.*
- *Use a subinterface on R6 for the ATM matching the VCI number and ensure the latest method of PVC configuration is used on this router. For R5 ATM, use the physical interface and legacy PVC configuration; after you have configured your Layer 2 information, you may then add the Layer 3 addresses.*
- *Do not rely on inverse ARP.*

R6 requires a point-to-point subinterface named ATM1/0.99 with the PVC details configured under the separate PVC; R5 requires the legacy style with the **map-list** to achieve the PVC connectivity in this back-to-back configuration. The **map-list, ip 10.99.99.1 atm-vc 1 broadcast**, and **protocol ip 10.99.99.2** commands ensure that **inverse-arp** is not relied upon.

You can use whichever encapsulation suits the three tasks in Section 1.3 as it has not been defined which type must be used.

If you have successfully configured all items as in Example 1-7 and Example 1-8, you have scored 3 points.

Example 1-7 *R6 ATM Configuration and Map Verification*

```
interface ATM1/0
 no ip address
 no atm ilmi-keepalive
!
interface ATM1/0.99 point-to-point
 ip address 10.99.99.1 255.255.255.248
pvc 0/99
  protocol ip 10.99.99.2  broadcast
  encapsulation aal5snap
R6#show atm map
Map list ATM1/0.99pvc1 : PERMANENT
ip 10.99.99.2 maps to VC 1, VPI 0, VCI 99, ATM1/0.99
        , broadcast
```

Example 1-8 *R6 ATM Configuration and Map Verification*

```
interface ATM3/0
 ip address 10.99.99.2 255.255.255.248
 map-group atm
 atm pvc 1 0 99 aal5snap
 no atm ilmi-keepalive
!
map-list atm
 ip 10.99.99.1 atm-vc 1 broadcast

R5#show atm map
Map list atm : PERMANENT
ip 10.99.99.1 maps to VC 1
        , broadcast
```

Section 2: IP IGP Protocols (28 Points)

Section 2.1: RIP (16 Points)

- *On all RIP router, ensure that **version 2** is used under the process.*

Add **version 2** under the RIP process. You receive no points here; this just ensures your routers behave correctly during the lab.

You should have at this point also enabled RIP for your networks using the **network** command and as a matter of good practice configured router interfaces that are not part of

the RIP domain as passive using the command **passive-interface** under the RIP process of each router.

- *Ensure that VLSM is supported on advertisements between all RIP routers.*

This is just a case of manually configuring the RIP routers to disable auto summarization mainly for the benefit of R3, which would otherwise receive a classfull network 10.0.0.0/8 route from R2. If you have configured this correctly as shown in Example 1-9 on all RIP routers with the resulting routing table shown for R3 in Example 1-10, you have scored 2 points.

Example 1-9 *RIP VLSM Configuration on R1, R2, R3, R4, and R6*

```
router rip
 no auto-summary
```

Example 1-10 *R3 RIP Routing Table Output*

```
R3#sh ip route

C    172.16.0.0/16 is directly connected, FastEthernet0/0
     10.0.0.0/8 is variably subnetted, 9 subnets, 3 masks
R       10.100.100.0/28 [120/2] via 172.16.0.1, 00:00:24, FastEthernet0/0
R       10.99.99.0/29 [120/4] via 172.16.0.1, 00:00:24, FastEthernet0/0
R       10.90.90.0/28 [120/1] via 172.16.0.1, 00:00:24, FastEthernet0/0
R       10.80.80.0/24 [120/2] via 172.16.0.1, 00:00:24, FastEthernet0/0
R       10.60.60.0/29 [120/4] via 172.16.0.1, 00:00:26, FastEthernet0/0
R       10.40.40.0/28 [120/3] via 172.16.0.1, 00:00:26, FastEthernet0/0
R       10.6.6.0/29 [120/4] via 172.16.0.1, 00:00:26, FastEthernet0/0
R       10.4.4.0/29 [120/3] via 172.16.0.1, 00:00:26, FastEthernet0/0
R       10.1.1.0/28 [120/2] via 172.16.0.1, 00:00:26, FastEthernet0/0
```

- *Add a loopback interface with the address of 60.60.60.1/24 onto R3 and advertise this out to R2 but ensure that it is not seen by the rest of your network; do not perform any configuration on R2 or R1.*

Add the loopback as Lo0 on R3 and enable the 60.60.60.0/24 network under RIP; this automatically advertises network 60.60.60.0/24 out to R2 and the rest of your RIP network over the 172.16.0.0/16 network, which you should have already configured according to the IGP diagram. The usual method to stop the propagation of this network would be via distribute lists, but the question states that R2 must see the network so you can not put a distribute list out on R3; the question also states that you can not configure R2 or R1 so you will have to configure R3. You need to get back to basics here and recall that RIP has a maximum hop count of 15 with 16 hops marked as unreachable so you will need to ensure that when network 60.60.60.0/24 egresses R3 the hop count is already set at 14. This way when R2 sees the route it knows that it has a hop count of 15 to reach it; it, in turn, will then advertise network 60.60.60.0/24 with a hop count of 16, which is, of course, unreachable and, hence, it will not be included in the routing table of R1 and beyond. To achieve the artificial hop count, an **offset-list** is required for network 60.60.60.0/24 on R3. If you have

configured this correctly as in Example 1-11 with validation shown in Example 1-12 and Example 1-13, you have scored 4 points.

You could have also gained full marks for advertising the loopback interface on R3 within RIP as a connected interface and assigned a metric of 15 to this route, which provides exactly the same result.

Example 1-11 *R3 Hop Count Configuration*

```
interface Loopback0
 ip address 60.60.60.1 255.255.255.0
!
router rip
 version 2
 offset-list 1 out 14 FastEthernet0/0
 network 60.0.0.0
 network 172.16.0.0
!
access-list 1 permit 60.60.60.0
```

Example 1-12 *R2 Routing Entry for 60.60.60.0/24*

```
R2#show ip route 60.60.60.0
Routing entry for 60.60.60.0/24
  Known via "rip", distance 120, metric 15
  Redistributing via rip
  Last update from 172.16.0.2 on FastEthernet0/0, 00:00:15 ago
  Routing Descriptor Blocks:
  * 172.16.0.2, from 172.16.0.2, 00:00:15 ago, via FastEthernet0/0
      Route metric is 15, traffic share count is 1
```

Example 1-13 *R1 RIP* **debug**

```
R1#debug ip rip
2w1d: RIP: received v2 update from 10.90.90.1 on Serial0/0
2w1d:      60.60.60.0 in 16 hops (inaccessible)
2w1d:      172.16.0.0 in 1 hops
```

- *Configure R3 to unicast its RIP routing updates to R2. Do not use the* **neighbor** *command to achieve this but consider using other IP features to aid you.*

Normally, you would use the **neighbor** command in conjunction with **passive-interface** to ensure that a router unicasts its routing updates instead of multicasting them in the usual manner. To achieve this without the neighbor command, you will need to use NAT to turn a multicast into a unicast; this is your additional IP feature. A simple NAT statement causing any packet with a destination address as a multicast to destination address 224.0.0.9 with the UDP port equal to that of RIP (520) to be converted into a destination address of 172.16.0.1 (R2 FastEthernet0/0) will cause R3 to now unicast its routing updates directly to R3.

If you have configured this correctly as in Example 1-14 and with the resulting output on R2 as shown in Example 1-15, you have scored 6 points.

Example 1-14 *R3 NAT Configuration and* **debug**

```
interface FastEthernet0/0
 ip address 172.16.0.2 255.255.0.0
 ip nat outside
!

ip nat outside source static udp 172.16.0.1 520 224.0.0.9 520

R3#debug ip nat det
IP NAT detailed debugging is on
R3#clear ip route *
R3#

00:57:29: NAT: i: udp (172.16.0.2, 520) -> (224.0.0.9, 520) [0]
00:57:29: NAT: s=172.16.0.2, d=224.0.0.9->172.16.0.1 [0]
```

Example 1-15 *R2 RIP* **debug**

```
R2#debug ip pack det
IP packet debugging is on (detailed)
R2#
00:54:56: IP: s=172.16.0.2 (FastEthernet0/0), d=172.16.0.1 (FastEthernet0/0), len 5
2, rcvd 3
00:54:56:     UDP src=520, dst=520
```

- *Ensure that VLAN2 is advertised to the RIP domain as a /28 network. Do not use either RIP or EIGRP features to accomplish this. You can, however, configure R6.*

VLAN2 has a subnet mask of /24, and as such, the RIP domain would see this as network 10.80.80.0/24.

You could quite easily summarize network 10.80.80.0/24 within RIP or later within EIGRP to change the network to 10.80.80.0/28, but the question clearly states that no RIP or EIGRP feature must be used. The lab rules are also not static routes; policy routing won't help as the network should be present in all routing tables so the only way to get VLAN2 from a /24 into a /28 is to think laterally and add a secondary address on R6 FastEthernet0/0 within the /28 range (i.e., 10.80.80.14/28). This will then ensure the network 10.80.80.0/28 is advertised into the RIP domain.

NOTE The new RIP advertisement of 10.80.80.0/28 will be received by R1, which already has a connected interface into the real 10.80.80.0/24 network. This is a longer match than its own connected interface and, hence, will cause suboptimal routing for R1 to communicate on VLAN2 within the range of the /24 subnet. A **distribute-list** must be used on R1 to filter this network. Remember that the RIP route for this network could arrive on both the Frame Relay interface and the BRI if the Frame Relay network fails later in the lab; as such the **distribute-list** is required in-bound on both interfaces.

If you have configured this correctly including filtering network 10.80.80.0/28 from entering R1 as in Example 1-16 and Example 1-17, you have scored 4 points. If you have only configured the **distribute-list** on the Frame Relay network, you have only scored 2 points.

Example 1-16 *R6 Secondary Address Configuration*

```
interface FastEthernet0/0
 ip address 10.80.80.14 255.255.255.240 secondary
 ip address 10.80.80.2 255.255.255.0
```

Example 1-17 *R1 RIP Distribute-List Configuration*

```
router rip
distribute-list 1 in Serial0/1.101
 distribute-list 1 in BRI0/0
!
access-list 1 deny   10.80.80.0 0.0.0.15
access-list 1 permit any
```

Section 2.2: EIGRP (5 Points)

You should have configured EIGRP using AS10 as shown in Figure 1-13 on R5, R6, R7, and R8. R6 has RIP enabled on the Frame Relay network, so you can either use a **network** statement for each EIGRP required interface with an inverse mask or simply use the **passive-interface** command as required. All EIGRP routers should also have auto summarization disabled using the command **no auto-summary**. No extra points here in Lab 1, but you will find in later labs that you will earn points for the correct basic configuration.

NOTE The IGP questions do not stipulate if R6 should advertise it's loopback interface via RIP or EIGRP because R6 runs both protocols, in this case it is prudent to do so in both instances.

- *R8 is very low on memory and CPU resource; accommodate this information within the configuration on R8.*

EIGRP supports stub routing, which improves network stability, reduces resource, and simplifies configuration. R8 does not participate in any summary advertisements so it

purely requires **eigrp stub connected** configured under its EIGRP process to ensure that its connected interfaces are successfully advertised out to its neighbors. If you have configured this correctly as in Example 1-18, you have scored 3 points.

Example 1-18 *R8 EIGRP Stub-Routing Configuration and R6 EIGRP Neighbor Output*

```
router eigrp 10
 network 10.0.0.0
 no auto-summary
 eigrp stub connected

R6#sh ip eigrp neighbors detail
IP-EIGRP neighbors for process 10
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq Type
                                            (sec)         (ms)       Cnt Num
2   10.99.99.2              Se0/0            167 05:30:02    4   200  0  2
   Version 12.2/1.2, Retrans: 6, Retries: 0
1   10.60.60.2              Fa0/1             12 05:30:02  340  2040  0  3
   Version 12.1/1.2, Retrans: 0, Retries: 0
0   10.80.80.3              Fa0/0             14 05:30:05    9   200  0  4
   Version 12.1/1.2, Retrans: 2, Retries: 0
   Stub Peer Advertising ( CONNECTED ) Routes
```

- *Configure R8 to have an EIGRP hello interval of 25 seconds on its FastEthernet0/0 interface.*

The EIGRP hello interval is by default set at 5 seconds for FastEthernet. This is not a difficult question but you must ensure if you are changing any EIGRP interval that you should also configure that of your neighbors on the common subnet exactly the same otherwise your neighbor adjacencies will be fluctuating as will your routing table. You should also be aware that the EIGRP hold interval should be three times that of the hello interval otherwise you will experience difficulties in maintaining your neighbor relationship. You should, therefore, configure the **ip hold-time eigrp** interval on R8 under the FastEthernet0/0 as 75 seconds. Configure R6 under its FastEthernet0/0 with the same configuration as R8 as it is a neighbor to R8 on VLAN2. If you have configured this correctly as shown in Example 1-19, you have scored 2 points.

Example 1-19 *R8 and R6 EIGRP Hello and Hold Interval Configuration*

```
interface FastEthernet0/0
 ip hello-interval eigrp 10 25
 ip hold-time eigrp 10 75
```

Section 2.3: Redistribution (7 Points)

- *Redistribute IGP protocols to ensure full IP visibility between all routers.*

You can see via the IGP diagram in Figure 1-13 that there will only be one redistribution point required, this being R6.

Mutual redistribution between RIP and EIGRP is required. Don't forget your default metrics under each process otherwise the different protocols will have no means of

applying relevant metrics to the routes you wish to advertise. If you have configured your redistribution correctly as shown in Example 1-20 and Example 1-21 and have full IP visibility of all networks, you have scored 4 points.

Example 1-20 *R6 EIGRP Redistribution to RIP Configuration*

```
router rip
 version 2
 redistribute eigrp 10
 passive-interface default
 no passive-interface Serial5/0.103
 network 10.0.0.0
 default-metric 3
 no auto-summary
```

Example 1-21 *R6 RIP Redistribution to EIGRP Configuration*

```
router eigrp 10
 redistribute rip
 passive-interface default
 no passive-interface FastEthernet0/0
 no passive-interface ATM1/0.99
 no passive-interface FastEthernet4/0
 network 10.0.0.0
 default-metric 100000 0 255 1 1500
 no auto-summary
```

- *As a safety precaution, ensure that R6 can not learn the EIGRP routes it previously advertised into the RIP domain back from R4.*

This question is just a straightforward practice of distribute lists and ensuring that the correct networks are filtered. In this scenario, R6 would ignore any routes back from RIP to which it had redistributed into RIP originally from EIGRP because of the external EIGRP route feature (any routes redistributed into EIGRP are subject to an increased Administritive Distance from 90 to 170). The redistributed RIP routes would simply be ignored. To answer the question as requested, though, you will need to configure a **distribute-list** within RIP on R6 Serial5/0.103, which blocks the EIGRP routes that R6 advertises out to the RIP domain. Do not include the connected interfaces on R6 in your ACL as these would be advertised within the RIP domain anyway and not redistributed into RIP from EIGRP. If you have configured this correctly as shown in Example 1-22, you have scored 3 points.

Example 1-22 *R6 Distribution List Configuration*

```
router rip
 distribute-list 1 in Serial5/0.103
!
access-list 1 deny   10.8.8.8
access-list 1 deny   10.5.5.4 0.0.0.3
access-list 1 deny   10.7.7.0 0.0.0.15
access-list 1 deny   10.50.50.0 0.0.0.7
access-list 1 permit any
```

Section 3: ISDN (8 Points)

- *Ensure that VLAN3 and R4 Lo0 are accessible from R1 and beyond should the Frame Relay network fail either physically or logically. If VLAN3 and R4 Lo0 networks are restored while the ISDN line is active, ensure that traffic is routed over the Frame Relay network to these destinations immediately.*

As no static routes are permitted and backing up the Frame Relay interface will not help as this only works if the Frame Relay interface is physically down, the only option will be to use the **dialer-watch** feature. Both networks must be down before the router dials out so VLAN3 and R4 Lo0 should be added to a **dialer watch-list** and corresponding **dialer watch-group** number under the BRI interface on R1. R1 is used to dial out as the question states that the two networks should be accessible from R1. You are also later advised that only R1 should dial into R4.

You should notice that when you fail the Frame Relay network to test this that after the ISDN is activated and the Frame Relay network is then restored that the routing table on R1 shows identical hop counts for all remote networks via R4 over both the Frame Relay and ISDN line as shown in Example 1-23.

This condition can keep the ISDN line from ever deactivating as the ISDN network can now be used as a valid means to transport data to the RIP advertised remote networks, you should also notice that the question requires that the Frame Relay routes should be used "immediately" when restored and, at this point, routers R1 and R4 can choose between Frame Relay and ISDN.

RIP obviously does not take into account the bandwidth of available routes. You, therefore, need to make the ISDN routes less desirable and add additional hop count to RIP using an **offset-list** on R4 and R1 out over the ISDN line (inbound over both routers will also be acceptable). This ensures when the Frame Relay is restored and for the period where both Frame Relay and ISDN lines are active and receiving RIP routes that the hop count is more favorable over Frame Relay because of the additional hop count incurred over ISDN after the **offset-list** is applied.

The ISDN line can not be used to route traffic while a higher-speed Frame Relay connection is available as shown in Example 1-25. **Dialer-watch** does not require interesting traffic to trigger the dial so the **dialer-list** should be an implicit deny of any IP traffic; otherwise, any traffic will potentially keep the line up after initiated. It is better practice and shows a better understanding of the dialer-watch process to, therefore, have the following **dialer-list** on R1; **dialer-list 10 protocol ip deny**. You will find with this strict policing of the interesting traffic, your ISDN line will stay down when the networks are restored over the Frame Relay. If you have configured this question correctly as in Example 1-24, you have scored 5 points; if you have used a **dialer-list** that denies RIP and the line stays down, you have only scored 3 points. Test your scenario thoroughly if you have first denied RIP then allowed all other IP traffic and also not applied the **offset-list**; you could find that with two routes in the routing table with identical metrics that traffic, such as BGP, will toggle between the two

routes and keep the line up constantly. In addition, other IP traffic could be classed as interesting and keep the line up.

Example 1-23 *R1 Routing Table Pre Offset-List with the ISDN Line Active After the Frame Relay Network Has Been Restored*

```
R1#sh ip route
R    172.16.0.0/16 [120/1] via 10.90.90.1, 00:00:22, Serial0/0
     10.0.0.0/8 is variably subnetted, 16 subnets, 5 masks
R       10.8.8.8/32 [120/4] via 10.10.10.2, 00:00:14, BRI0/0
                    [120/4] via 10.100.100.2, 00:00:14, Serial0/1.101
C       10.10.10.2/32 is directly connected, BRI0/0
C       10.100.100.0/28 is directly connected, Serial0/1.101
R       10.99.99.0/29 [120/2] via 10.10.10.2, 00:00:14, BRI0/0
                      [120/2] via 10.100.100.2, 00:00:14, Serial0/1.101
R       10.60.60.0/29 [120/2] via 10.10.10.2, 00:00:14, BRI0/0
                      [120/2] via 10.100.100.2, 00:00:14, Serial0/1.101
R       10.50.50.0/29 [120/4] via 10.10.10.2, 00:00:14, BRI0/0
                      [120/4] via 10.100.100.2, 00:00:14, Serial0/1.101
R       10.40.40.0/28 [120/1] via 10.10.10.2, 00:00:16, BRI0/0
                      [120/1] via 10.100.100.3, 00:00:16, Serial0/1.101
R       10.7.7.0/28 [120/4] via 10.10.10.2, 00:00:16, BRI0/0
                    [120/4] via 10.100.100.2, 00:00:16, Serial0/1.101
R       10.6.6.0/29 [120/2] via 10.10.10.2, 00:00:16, BRI0/0
                    [120/2] via 10.100.100.2, 00:00:16, Serial0/1.101
R       10.4.4.0/29 [120/1] via 10.10.10.2, 00:00:16, BRI0/0
                    [120/1] via 10.100.100.3, 00:00:16, Serial0/1.101
C       10.80.80.0/24 is directly connected, FastEthernet0/0
C       10.90.90.0/28 is directly connected, Serial0/0
C       10.1.1.0/28 is directly connected, Loopback0
C       10.10.10.0/28 is directly connected, BRI0/0
C       10.90.90.1/32 is directly connected, Serial0/0
R       10.5.5.4/30 [120/4] via 10.10.10.2, 00:00:16, BRI0/0
                    [120/4] via 10.100.100.2, 00:00:16, Serial0/1.101
R1#sh isdn history
--------------------------------------------------------------------------------
                               ISDN CALL HISTORY
--------------------------------------------------------------------------------
Call History contains all active calls, and a maximum of 100 inactive calls.
Inactive call data will be retained for a maximum of 15 minutes.
--------------------------------------------------------------------------------
Call    Calling      Called       Remote  Seconds Seconds Seconds Charges
Type    Number       Number       Name    Used    Left    Idle    Units/Currency
--------------------------------------------------------------------------------
Out                    2222                   82     37       82      0
--------------------------------------------------------------------------------
```

NOTE The routing table output is taken after a Frame Relay failure is restored and the ISDN line is still active. The shading shows you the two available routes with the identical hop count on R1 before the **offset-list** is applied.

Example 1-24 *Increasing the Hop Count Out of R4 and R1 ISDN Configuration*

```
R1
router rip
offset-list 0 out 2 BRI0/0
R4
router rip
 offset-list 0 out 2 Dialer0
```

NOTE **offset-list 0** will apply the chosen additional hop count (2) to all networks being advertised from R4 and R1 out of their interfaces BRI0/0. A similar configuration could be placed on each BRI0/0 but inbound.

Example 1-25 *R1 Routing Table Post Offset-List with the ISDN Line Active After the Frame Relay Network Has Been Restored*

```
R1#sh ip route
R    172.16.0.0/16 [120/1] via 10.90.90.1, 00:00:09, Serial0/0
     10.0.0.0/8 is variably subnetted, 16 subnets, 5 masks
R       10.8.8.8/32 [120/4] via 10.100.100.2, 00:00:28, Serial0/1.101
C       10.10.10.2/32 is directly connected, BRI0/0
C       10.100.100.0/28 is directly connected, Serial0/1.101
R       10.99.99.0/29 [120/2] via 10.100.100.2, 00:00:28, Serial0/1.101
R       10.60.60.0/29 [120/2] via 10.100.100.2, 00:00:28, Serial0/1.101
R       10.50.50.0/29 [120/4] via 10.100.100.2, 00:00:28, Serial0/1.101
R       10.40.40.0/28 [120/1] via 10.100.100.3, 00:00:28, Serial0/1.101
R       10.7.7.0/28 [120/4] via 10.100.100.2, 00:00:28, Serial0/1.101
R       10.6.6.0/29 [120/2] via 10.100.100.2, 00:00:28, Serial0/1.101
R       10.4.4.0/29 [120/1] via 10.100.100.3, 00:00:28, Serial0/1.101
C       10.80.80.0/24 is directly connected, FastEthernet0/0
C       10.90.90.0/28 is directly connected, Serial0/0
C       10.1.1.0/28 is directly connected, Loopback0
C       10.10.10.0/28 is directly connected, BRI0/0
C       10.90.90.1/32 is directly connected, Serial0/0
R       10.5.5.4/30 [120/4] via 10.100.100.2, 00:00:00, Serial0/1.101

R1#sh isdn hist
--------------------------------------------------------------------------------
                                ISDN CALL HISTORY
--------------------------------------------------------------------------------
Call History contains all active calls, and a maximum of 100 inactive calls.
```

continues

Example 1-25 *R1 Routing Table Post Offset-List with the ISDN Line Active After the Frame Relay Network Has Been Restored (Continued)*

```
Inactive call data will be retained for a maximum of 15 minutes.
-------------------------------------------------------------------------------
Call    Calling      Called       Remote  Seconds Seconds Seconds Charges
Type    Number       Number       Name    Used    Left    Idle    Units/Currency
-------------------------------------------------------------------------------
Out                    2222                    92     27       92      0
-------------------------------------------------------------------------------
```

NOTE The routing table output is taken after a Frame Relay failure is restored and the ISDN line is still active. This shows that the ISDN routes are no longer entered into the routing table on R1 because of the increased hop count over this environment. The routing table on R4 will act in exactly the same manner.

- *Configure R1 so that if half of the ISDN traffic to R4 is of an unacceptable quality, the line is automatically disconnected.*

Configure **ppp quality 50** under both R1 and R4 BRI0/0 interfaces, the figure (percentage) is for both incoming and outgoing directions on the interface, PPP will drop the line if the quality falls below 50 percent and initiate a timer before re-establishing the link. If you have configured this correctly, you have scored 1 point.

- *Allow only R1 to dial into R4. Do not use any PPP feature in your solution.*

The question is not seeking configuration of CHAP on both routers as any router configured with the correct CHAP password could emulate R1 and gain access to R4. It is, therefore, required to configure R4 with **isdn caller 1111** if using legacy DDR or **dialer-caller 1111** if using dialer profiles to ensure that only R1, which is connected to the ISDN number 1111, can actually gain access by having R4 check the CLI before answering. You may have automatically assumed this must require CHAP but there is not sufficient detail in the question to suggest that CHAP or PAP is required. These are both also PPP features so it is disallowed anyway. If you have configured this correctly, you have scored 2 points.

NOTE Your ISDN line or simulator must support CLI to test this feature.

- *Do not allow the ISDN to flap if the Frame Relay network goes up and down; only allow the line to be dropped if the Frame Relay connectivity is deemed to be reliable for 90 seconds.*

By default, the ISDN line will be dropped when dialer-watch again has visibility if the networks listed in the dialer watch-list. To ensure the line remains active for 90 seconds the command **dialer watch-disable 90** should be added to the BRI0/0 interface of R1. If you have configured this correctly, you have scored 1 point.

Example 1-26 and Example1-27 show the full final ISDN and relevant RIP configuration required for the ISDN backup on R1 and R4, using a mix of legacy and dialer profile commands.

Example 1-26 *R1 Final ISDN and Relevant RIP Configuration*

```
interface BRI0/0
 ip address 10.10.10.1 255.255.255.240
 encapsulation ppp
 dialer watch-disable 90
 dialer string 2222
 dialer watch-group 5
 dialer-group 10
 isdn switch-type basic-net3
 no peer neighbor-route
 ppp quality 50
!
router rip
 version 2
 passive-interface default
 no passive-interface BRI0/0
 offset-list 0 out 2 BRI0/0
 network 10.0.0.0
!
dialer watch-list 5 ip 10.4.4.0 255.255.255.240
dialer watch-list 5 ip 10.40.40.0 255.255.255.240
dialer-list 10 protocol ip deny
```

Example 1-27 *R4 Final ISDN and Relevant RIP Configuration*

```
interface BRI0/0
encapsulation ppp
isdn switch-type basic-net3
dialer pool-member 1

!
interface Dialer0
 ip address 10.10.10.2 255.255.255.240
 encapsulation ppp
 dialer pool 1
 dialer-group 10
 ppp quality 50
 dialer-caller 1111
!
router rip
 version 2
 passive-interface default
 no passive-interface Serial0/0.1
 no passive-interface Dialer0
 offset-list 0 out 2 D0
 dialer-list 10 protocol ip permit
```

Section 4: EGP Protocols (17 Points)

- *Configure BGP as shown in Figure 1-14 with the following peering: R3–R2, R8–R2, R6–R2, R6–R8, R7–R8, R6–R5. Ensure that most suitable interfaces are used to maintain resilience for BGP peering (except for R2 and R3 that use 172.16.0.0/16 addresses for all peering to and from these routers).*

You are required to configure the peering between the BGP autonomous systems as described. You should ensure that **no synchronization** is configured on all IBGP routers (R2, R3, R6, R7, and R8) as BGP in this scenario is not synchronized with the underlying IGP and, hence, it would not be able to advertise transit routes to external autonomous sytems. As requested, you should peer from your loopback interfaces where present to maintain resiliency except for R2 and R3. This requires BGP Multihop on all external BGP connections sourced from the loopbacks because, by default, a BGP speaker drops any UPDATE message from its EBGP peer, unless it is on the same connected network. By adding a number of hops to the command, you can ensure that the peering is achieved regardless of the traffic path taken (**ebgp-multihop 5**). Multihop should be used in conjunction with the **update-source** command to ensure that peering is maintained correctly by making the source IP address used for the BGP session the same as the remote BGP speakers neighbor statement address and not that of the connected interface. If you have configured this correctly as in Example 1-28 through Example 1-32, you have scored 3 points.

Example 1-28 *R2 Initial BGP Peering Configuration*

```
router bgp 10
 no synchronization
 neighbor 10.6.6.6 remote-as 100
 neighbor 10.6.6.6 ebgp-multihop 5
 neighbor 10.6.6.6 update-source FastEthernet0/0
 neighbor 10.8.8.8 remote-as 100
 neighbor 10.8.8.8 ebgp-multihop 5
 neighbor 10.8.8.8 update-source FastEthernet0/0
 neighbor 172.16.0.2 remote-as 10
```

Example 1-29 *R3 Initial BGP Peering Configuration*

```
router bgp 10
 no synchronization
 neighbor 172.16.0.1 remote-as 10
```

Example 1-30 *R5 Initial BGP Peering Configuration*

```
router bgp 1000
 neighbor 10.6.6.6 remote-as 100
 neighbor 10.6.6.6 ebgp-multihop 5
 neighbor 10.6.6.6 update-source Loopback0
```

Example 1-31 *R6 Initial BGP Peering Configuration*

```
router bgp 100
 no synchronization
 neighbor 10.8.8.8 remote-as 100
 neighbor 10.8.8.8 update-source Loopback0
 neighbor 10.5.5.5 remote-as 1000
 neighbor 10.5.5.5 update-source Loopback0
 neighbor 10.5.5.5 ebg-multihop 5
 neighbor 172.16.0.1 remote-as 10
 neighbor 172.16.0.1 ebgp-multihop 5
 neighbor 172.16.0.1 update-source Loopback0
```

NOTE There is no **ebg-multihop** required to peer to 10.8.8.8 in AS100 as this is internal BGP (IBGP) and not external BGP (EBGP).

Example 1-32 *R7 Initial BGP Peering Configuration*

```
router bgp 100
 no synchronization
 neighbor 10.8.8.8 remote-as 100
 neighbor 10.8.8.8 update-source Loopback0
```

NOTE R8 initial BGP peering configuration is covered in the following question.

- *Ensure minimal configuration on R8.*

R8 peers to three other routers, two of which belong to the same AS. You can, therefore, take advantage of BGP peer groups to reduce the required configuration for the policies to R6 and R7. You should be aware that full IBGP peering between R8-R7-R6 does not exist. Both R7 and R6 peer to R8 so, as well as running peer-groups, R8 should also be a route-reflector to overcome the IBGP peering problem. If you have configured this correctly as in Example 1-33, you have scored 1 point.

Example 1-33 *R8 Initial BGP Peering Configuration*

```
router bgp 100
 no synchronization
 neighbor cisco peer-group
 neighbor cisco remote-as 100
 neighbor cisco update-source Loopback0
 neighbor cisco route-reflector-client
 neighbor 10.6.6.6 peer-group cisco
 neighbor 10.7.7.7 peer-group cisco
```

continues

Example 1-33 *R8 Initial BGP Peering Configuration (Continued)*

```
 neighbor 172.16.0.1 remote-as 10
 neighbor 172.16.0.1 ebgp-multihop 5
 neighbor 172.16.0.1 update-source Loopback0
```

- *Inject the following networks into BGP via new loopback interfaces:*

  R3: 20.200.200.1/24 and 20.20.20.1/24

  R5: 20.20.20.1/24 and 200.20.20.1/24

  R7: 30.30.30.30/29

You should add the loopback interface and address as requested. The loopbacks are advertised into BGP simply with the network command. You should notice that both R3 and R5 will be advertising the same network (more of this later). If you have configured this correctly as in Example 1-34 through Example 1-36, you have scored 1 point.

Example 1-34 *R3 Loopback and BGP Advertisement Configuration*

```
interface Loopback1
 ip address 20.20.20.1 255.255.255.0
!
interface Loopback2
 ip address 20.200.200.1 255.255.255.0
!
router bgp 10
 network 20.20.20.0 mask 255.255.255.0
 network 20.200.200.0 mask 255.255.255.0
```

Example 1-35 *R5 Loopback and BGP Advertisement Configuration*

```
interface Loopback1
 ip address 20.20.20.1 255.255.255.0
!
interface Loopback2
 ip address 200.20.20.1 255.255.255.0
!
router bgp 1000
 network 20.20.20.0 mask 255.255.255.0
 network 200.20.20.0
```

NOTE Network 200.20.20.0 on R5 does not require an explicit **mask** because of being a class C network. As such, it will automatically summarize on the classfull network boundary if the **mask** command is omitted.

Example 1-36 *R7 Loopback and BGP Advertisement Configuration*

```
interface Loopback1
 ip address 30.30.30.30 255.255.255.248
!
 router bgp 100
  network 30.30.30.24 mask 255.255.255.248
```

Example 1-37 shows a snapshot of the BGP routing tables for all BGP routers at this point in time. You can use this as a quick check to ensure you see the advertised networks correctly on all routers and specifically on R6 before the complex BGP scenarios begin.

Example 1-37 **show ip bgp** *Output from Each BGP Router*

```
R2#sh ip bgp
BGP table version is 24, local router ID is 172.16.0.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*>i20.20.20.0/24    172.16.0.2               0    100      0 i
*>i20.200.200.0/24  172.16.0.2               0    100      0 i
*> 30.30.30.24/29   10.6.6.6                              0 100 i
*                   10.8.8.8                              0 100 i
*  200.20.20.0      10.8.8.8                              0 100 1000 i
*>                  10.6.6.6                              0 100 1000 i
R3#sh ip bgp
BGP table version is 20, local router ID is 20.200.200.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 20.20.20.0/24    0.0.0.0                  0         32768 i
*> 20.200.200.0/24  0.0.0.0                  0         32768 i
*>i30.30.30.24/29   10.6.6.6                      100      0 100 i
*>i200.20.20.0      10.6.6.6                      100      0 100 1000 i
R5#sh ip bgp
BGP table version is 13, local router ID is 200.20.20.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*  20.20.20.0/24    10.6.6.6                              0 100 10 i
*>                  0.0.0.0                  0         32768 i
*> 20.200.200.0/24  10.6.6.6                              0 100 10 i
*> 30.30.30.24/29   10.6.6.6                              0 100 i
*> 200.20.20.0      0.0.0.0                  0         32768 i
R6#sh ip bgp
BGP table version is 5, local router ID is 10.6.6.6
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete
```

continues

Example 1-37 **show ip bgp** *Output from Each BGP Router (Continued)*

```
   Network          Next Hop            Metric LocPrf Weight Path
*> 20.20.20.0/24    172.16.0.1                             0 10 i
*                   10.5.5.5                 0             0 1000 i
* i                 172.16.0.1                    100      0 10 i
*> 20.200.200.0/24  172.16.0.1                             0 10 i
* i                 172.16.0.1                    100      0 10 i
*>i30.30.30.24/29   10.7.7.7                 0    100      0 i
*> 200.20.20.0      10.5.5.5                 0             0 1000 i
R7#sh ip bgp
BGP table version is 18, local router ID is 30.30.30.30
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*>i20.20.20.0/24    172.16.0.1                    100      0 10 i
*>i20.200.200.0/24  172.16.0.1                    100      0 10 i
*> 30.30.30.24/29   0.0.0.0                  0         32768 i
*>i200.20.20.0      10.5.5.5                 0    100      0 1000 i
R8#sh ip bgp
BGP table version is 18, local router ID is 10.8.8.8
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
* i20.20.20.0/24    172.16.0.1                    100      0 10 i
*>                  172.16.0.1                             0 10 i
* i20.200.200.0/24  172.16.0.1                    100      0 10 i
*>                  172.16.0.1                             0 10 i
*>i30.30.30.24/29   10.7.7.7                 0    100      0 i
*>i200.20.20.0      10.5.5.5                 0    100      0 1000 i
```

- *Ensure that R6 BGP routing table prefers to use AS1000 for network 20.20.20.0/24; do not use BGP weight, BGP local preference, MED, neighbor metric related statements, metric manipulation, summarization, or prepending to achieve this. Perform configuration on R6 only.*

When you look at the BGP routing table on R6 as in Example 1-37, you can see that it has a best path to 20.20.20.0/24 from AS10 next hop 172.16.0.1 (R2 propagating the route from R3). It would be very straightforward to manipulate weight or local preference to ensure R6 prefers the same route received from AS1000 (R5), but the question is very strict. You can tell that both routes are EBGP and, hence, the distance to both routes is an Administrative Distance of 20, so why is the route to 172.16.0.1 preferred?

Example 1-37 also shows the routing table on R6. You can see that the IGP metric to neighbor (R2) 172.16.0.1 is 3 compared to that of 208384 to neighbor (R5) 10.5.5.5 and this is why the next hop to network 20.20.20.0/24 on R6 is 172.16.0.1. This is in accordance to Step 8 (prefer the path with the lowest IGP metric to the BGP next hop) in the 13 steps of Best Path Algorithm according to Cisco. You could reduce the IGP metric to R5 down to 0 on R6 by peering directly to the connected ATM interface (10.99.99.2) on R5 from R6 instead of the loopback on R5 from R6. This would ensure that R6 then prefers the route to

network 20.20.20.0/24 from R5, but this would infringe how you have been asked to peer in the original BGP setup question. As you can not manipulate favored attributes such as weight, local preference, AS-Path, summarization, or metrics, you are only left with Step 5 (prefer the path with the lowest origin type: IGP is lower than EGP, and EGP is lower than INCOMPLETE). As can be seen in Example 1-38, all routes for network 20.20.20.0/24 have an origin of IGP. As you may only configure R6, you can place an inbound **route-map** for neighbor 172.16.0.1 and change the origin of the received route for network 20.20.20.0/24 on R6. The solution as shown in Example 1-39 changes the origin to "incomplete" but if you have configured it to "external," this is also acceptable. If you have configured this correctly with the resulting BGP routing table on R6 as shown in Example 1-40 and BGP show output for network 20.20.20.0/24 as shown in Example 1-41, you have scored 7 points.

Example 1-38 **show ip route** *and* **show ip bgp** *Output*

```
R6#sh ip route
B    200.20.20.0/24 [20/0] via 10.5.5.5, 00:00:38
     20.0.0.0/24 is subnetted, 2 subnets
B       20.200.200.0 [20/0] via 172.16.0.1, 00:00:38
B       20.20.20.0 [20/0] via 172.16.0.1, 00:00:38
R    172.16.0.0/16 [120/3] via 10.100.100.1, 00:00:02, Serial5/0.103
     10.0.0.0/8 is variably subnetted, 16 subnets, 5 masks
D       10.8.8.8/32 [90/156160] via 10.80.80.3, 00:20:29, FastEthernet0/0
R       10.90.90.0/28 [120/2] via 10.100.100.1, 00:00:02, Serial5/0.103
R       10.1.1.0/28 [120/2] via 10.100.100.1, 00:00:03, Serial5/0.103
D       10.7.7.0/28 [90/156160] via 10.60.60.2, 00:20:30, FastEthernet4/0
R       10.40.40.0/28 [120/1] via 10.100.100.3, 00:00:03, Serial5/0.103
R       10.10.10.0/28 [120/1] via 10.100.100.3, 00:00:03, Serial5/0.103
R       10.4.4.0/29 [120/1] via 10.100.100.3, 00:00:03, Serial5/0.103
C       10.100.100.0/28 is directly connected, Serial5/0.103
D       10.50.50.0/29 [90/82944] via 10.99.99.2, 00:21:08, ATM1/0.99
C       10.99.99.0/29 is directly connected, ATM1/0.99
C       10.6.6.0/29 is directly connected, Loopback0
C       10.60.60.0/29 is directly connected, FastEthernet4/0
C       10.80.80.0/28 is directly connected, FastEthernet0/0
C       10.80.80.0/24 is directly connected, FastEthernet0/0
R       10.90.90.1/32 [120/2] via 10.100.100.1, 00:00:04, Serial5/0.103
D       10.5.5.4/30 [90/208384] via 10.99.99.2, 00:21:08, ATM1/0.99
     30.0.0.0/29 is subnetted, 1 subnets
B       30.30.30.24 [200/0] via 10.7.7.7, 00:00:40

R6#sh ip bgp
BGP table version is 7, local router ID is 10.6.6.6
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 20.20.20.0/24    172.16.0.1                             0 10 i
*                   10.5.5.5                 0             0 1000 i
* i                 172.16.0.1                    100      0 10 i
*> 20.200.200.0/24  172.16.0.1                             0 10 i
* i                 172.16.0.1                    100      0 10 i
*>i30.30.30.24/29   10.7.7.7                  0    100      0 i
*> 200.20.20.0      10.5.5.5                 0             0 1000 i
```

Example 1-39 *R6* **show ip bgp 20.20.20.0** *Output*

```
R6#sh ip bgp 20.20.20.0
BGP routing table entry for 20.20.20.0/24, version 6
Paths: (3 available, best #1, table Default-IP-Routing-Table)
  Advertised to non peer-group peers:
  10.5.5.5 10.8.8.8
  10
    172.16.0.1 (metric 3) from 172.16.0.1 (172.16.0.1)
      Origin IGP, localpref 100, valid, external, best
  1000
    10.5.5.5 (metric 208384) from 10.5.5.5 (200.20.20.1)
      Origin IGP, metric 0, localpref 100, valid, external
  10
    172.16.0.1 (metric 3) from 10.8.8.8 (10.8.8.8)
      Origin IGP, localpref 100, valid, internal
```

Example 1-40 *R6 Origin Configuration*

```
router bgp 100
neighbor 172.16.0.1 route-map 20.20.20.0 in
!
access-list 2 permit 20.20.20.0
route-map 20.20.20.0 permit 10
 match ip address 2
 set origin incomplete
!
route-map 20.20.20.0 permit 10
```

Example 1-41 *R6* **show ip bgp** *Output*

```
R6#sh ip bgp
BGP table version is 5, local router ID is 10.6.6.6
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 20.20.20.0/24    10.5.5.5                 0             0 1000 i
*                   172.16.0.1                             0 10 ?
* i                 172.16.0.1                    100      0 10 i
*> 20.200.200.0/24  172.16.0.1                             0 10 i
* i                 172.16.0.1                    100      0 10 i
*>i30.30.30.24/29   10.7.7.7                  0    100      0 i
*> 200.20.20.0      10.5.5.5                  0             0 1000 i
```

Example 1-42 *R6* **show ip bgp 20.20.20.0** *Output*

```
R6#sh ip bgp 20.20.20.0
BGP routing table entry for 20.20.20.0/24, version 2
Paths: (3 available, best #1, table Default-IP-Routing-Table)
  Advertised to non peer-group peers:
  10.8.8.8 172.16.0.1
  1000
    10.5.5.5 (metric 208384) from 10.5.5.5 (10.5.5.5)
```

Example 1-42 *R6* **show ip bgp 20.20.20.0** *Output (Continued)*

```
      Origin IGP, metric 0, localpref 100, valid, external, best
  10
    172.16.0.1 (metric 3) from 172.16.0.1 (172.16.0.1)
     Origin incomplete, localpref 100, valid, external
  10
    172.16.0.1 (metric 3) from 10.8.8.8 (10.8.8.8)
      Origin IGP, localpref 100, valid, internal
```

- *All BGP speakers are to be able to communicate with all advertised BGP networks.*

The BGP routes are in the BGP speakers routing tables so you should be able to ping the BGP networks, but can you? Without further configuration the answer is no. It should be painfully obvious that not all your routers are running BGP. You have not been requested to redistribute BGP into your IGP, so R1 and R4 will have no knowledge of any BGP networks. Example 1-43 shows what happens if you attempt to ping 30.30.30.30 from R2.

Example 1-43 *R2 Show IP Route Output and Connectivity Testing*

```
R2#sh ip route 30.30.30.30
Routing entry for 30.30.30.24/29
  Known via "bgp 10", distance 20, metric 0
  Tag 100, type external
  Last update from 10.6.6.6 00:28:44 ago
  Routing Descriptor Blocks:
  * 10.6.6.6, from 10.6.6.6, 00:28:44 ago
      Route metric is 0, traffic share count is 1
      AS Hops 1
R2#ping 30.30.30.30
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 30.30.30.30, timeout is 2 seconds:
U.U.U
Success rate is 0 percent (0/5)
```

As you can see in Example 1-43, R2 obviously has a route to the destination address 30.30.30.30 but R2 is receiving ICMP unreachable messages from R1 as R1 does not have visibility of network 30.30.30.24/29 and, hence, is dropping the packets and informing R2. Example 1-44 shows R1 and R4 have no visibility of the BGP networks.

Example 1-44 *R1 and R4 IGP out of Sync* **show output** *and R2 Traceroute*

```
R1#sh ip route 30.30.30.30
% Network not in table
R4#sh ip route 30.30.30.30
% Network not in table
R2#traceroute 30.30.30.30

Type escape sequence to abort.
Tracing the route to 30.30.30.30

  1 10.90.90.2 4 msec 4 msec 4 msec
  2 10.90.90.2 !H  *  !H
```

You, therefore, need to policy route at R1 and R4 for the specific BGP routes at these routers as shown in Example 1-45 and Example 1-46.

Example 1-45 *R1 Required Policy-Routing Configuration*

```
interface Serial0/0
 ip policy route-map as100-1000
!
interface Serial0/1.101 point-to-point
ip policy route-map as10
!
route-map as100-1000 permit 10
 match ip address 102
 set ip next-hop 10.100.100.2
!
route-map as10 permit 10
 match ip address 101
 set ip next-hop 10.90.90.1
!
access-list 101 permit ip any 20.200.200.0 0.0.0.255
access-list 102 permit ip any 30.30.30.24 0.0.0.7
access-list 102 permit ip any 200.20.20.0 0.0.0.255
```

NOTE Two separate route-maps are shown in Example 1-45 for traffic flowing towards network 20.200.200.0, which should be forwarded onto R2 and traffic flowing toward 30.30.30.0 and 200.20.20.0, which should be forwarded onto R6. No requirement exists to policy route traffic across the BRI on R1 or R4 as ISDN should only be used for the requirements stated in the questions.

No requirement exists to policy route traffic destined for network 20.20.20.0/24 on R1 as traffic will stay local to AS10 between R2 and R3 and never attempt to flow through R1. Similarly, R1 will never receive traffic destined for network 20.20.20.0/24 from AS100 as the preferred route will be to AS1000 R5 via R6.

Example 1-46 *R4 Required Policy-Routing Configuration*

```
interface Serial0/0.1 multipoint
ip policy route-map as10
!
route-map as10 permit 10
 match ip address 101
 set ip next-hop 10.100.100.1
!
route-map as10 permit 20
 match ip address 102
 set ip next-hop 10.100.100.2
!
access-list 101 permit ip any 20.200.200.0 0.0.0.255
access-list 102 permit ip any 200.20.20.0 0.0.0.255
access-list 102 permit ip any 30.30.30.24 0.0.0.7
```

NOTE Example 1-46 shows two separate route-map sequences for traffic flowing toward network 20.200.200.0, which should be forwarded onto R2, and traffic flowing toward 30.30.30.0 and 200.20.20.0, which should be forwarded onto R6.

No requirement exists to policy route traffic destined for network 20.20.20.0/24 on R4 as R6 will never send traffic to R4 destined for this network as it has a more preferable route to R5.

You should notice that both R7 and R8 show the next hop to 20.20.20.0/24 as 172.16.0.1. Traffic would egress R7 and R8 toward R2 172.16.0.1 and flow through R6 en route to R2, R6 would then send this to R5 as it is its own preferred route to 20.20.20.0/24. You therefore could consider policy routing at R6 for traffic sourced from R7 and R8 toward network 20.20.20.0/24. If you have configured this, it was a prudent action but the question did not ask you to accomplish this so no extra points or more, importantly, none deducted. You should seek advice from the proctor if an issue like this arises in your real exam, though. If you have configured this correctly as in Example 1-45 and Example 1-46, you have scored 5 points.

Section 5: Voice (6 Points)

- *Both phones should be able to ring each other using the numbers supplied. Use the most efficient method of transporting the voice from R1 to R4.*

Did you jump straight in with Voice over IP (VoIP)? VoIP is not the most efficient means of transporting voice between R1 and R4 as they have a dedicated Frame Relay connection between them. If you were to use VoIP, the voice would have to be encapsulated into IP and then into Frame Relay before it is even transmitted. The most efficient method is therefore to encapsulate directly into Frame Relay using Voice over Frame Relay (VoFR). VoFR requires a **map-class** on the Frame Relay DLCI. Otherwise, it will break the existing data connectivity between R1 and R4, so pay particular attention to the commands, which must include **frame-relay fragment** and under the physical Frame Relay interface, you must configure the command **frame-relay traffic shaping**. The **dial-peers** for the remote site numbers simply point to the DLCIs between R1 and R4. Ensure you configure the command **frame-relay voice bandwidth** under the Frame Relay **map-class** on routers R1 and R4; otherwise, your voice will not work. This value could be calculated exactly using the standard voice codec g729r8 with associated overhead and the number of required calls but the question does not request this. If you configured this correctly as in Example 1-47

and Example 1-48, you have scored 3 points. If you succeeded in the VoFR but have broken the data connectivity, you have scored no points.

Example 1-47 *R1 Voice and VoFR Configuration*

```
interface Serial0/1
 frame-relay traffic-shaping
!
interface Serial0/1.101 point-to-point
 frame-relay  interface-dlci 101
  class ccie
  vofr cisco
!
map-class frame-relay ccie
 frame-relay fair-queue
 frame-relay voice bandwidth 64000
 frame-relay fragment
!
dial-peer voice 1 pots
 destination-pattern 01256
 port 1/0/0
!
dial-peer voice 2 vofr
 destination-pattern 01189
 session target Serial0/1 101
```

Example 1-48 *R4 Voice and VoFR Configuration*

```
interface Serial0/0
  frame-relay traffic-shaping
!
interface Serial0/0.1 multipoint
frame-relay interface-dlci 100
  class ccie
  vofr cisco
 no frame-relay inverse-arp
!
map-class frame-relay ccie
frame-relay fair-queue
 frame-relay voice bandwidth 64000
frame-relay fragment
!
dial-peer voice 1 pots
 destination-pattern 01189
 port 1/1/0
!
dial-peer voice 2 vofr
 destination-pattern 01256
 session target Serial0/0 100
```

- *Ensure that voice is still available if the main connection between R1 and R4 fails.*

It is now time to configure VoIP, because if the Frame Relay network fails, you will need to run voice over the ISDN network between R1 and R4. The only method available to you is VoIP. Additional dial-peers are required each end pointing to the loopback IP addresses of each remote router. You must ensure that the VoFR is used before the VoIP so you will need to allocate a preference to the dial-peers. The lowest number dial-peers will have preference so it is advised that the VoFR dial-peers should be configured before the VoIP. Alternatively the VoIP dial-peers can have a priority manually configured higher than the default 0, which ensures if your dial-peer numbering is out of sync, the router still chooses VoFR before VoIP. Both methods of priority have been shown in the configuration for clarity. If you have configured this correctly as in Example 1-49 and Example 1-50, you have scored 2 points.

Example 1-49 *R1 VoIP Configuration*

```
dial-peer voice 3 voip
 preference 5
 destination-pattern 01189
 session target ipv4:10.4.4.4
```

Example 1-50 *R4 VoIP Configuration*

```
dial-peer voice 3 voip
 preference 5
 destination-pattern 01256
 session target ipv4:10.1.1.1
```

- *Make the phone on R1 also answer calls if 01962 is dialed on the R4 connected handsets; do not use number expansion to achieve this.*

You are required to simply add additional dial-peers for 01962 on R4 (both VoFR and VoIP) pointing to R1, and then configure R1 with an additional dial-peer POTS pointing to the original FXS phone port, which currently contains 01256. If you have configured this correctly as in Example 1-51 and Example 1-52, you have scored 1 point.

Example 1-51 *R1 01962 Configuration*

```
dial-peer voice 6 pots
 destination-pattern 01962
 port 1/0/0
```

Example 1-52 *R4 01962 Configuration*

```
dial-peer voice 4 VoFR
 destination-pattern 01962
 session target Serial0/0 100
!
dial-peer voice 5 voip
 preference 5
 destination-pattern 01962
 session target ipv4:10.1.1.1
```

Section 6: DLSw+ (4 Points)

- *Configure DLSw+ between VLAN2 and VLAN4; use routers R8 and R5. Peer from Lo0 on R8 and the VLAN4 interface on R5; ensure that R8 can accept DLSw+ connections from only unknown TCP peers.*

Peer as requested from Lo0 on R8 and configure this peer as promiscuous with TCP for the future connections. Configure R5 as requested and configure your bridging parameters to ensure the required connectivity. You should be concerned about this question; it is too easy and this should be ringing alarm bells. If you have configured this correctly as in Example 1-53 and Example 1-54, you have scored 2 points.

Example 1-53 *R5 Iinitial DLSw+ Configuration and* **show output**

```
dlsw local-peer peer-id 10.50.50.1
dlsw remote-peer 0 tcp 10.8.8.8
dlsw bridge-group 1
!
interface FastEthernet0/0
bridge-group 1
!
bridge 1 protocol ieee

R5#sh dlsw peer
Peers:                state     pkts_rx   pkts_tx  type  drops ckts TCP   uptime
 TCP 10.8.8.8        CONNECT      11914     11344  conf      0    0   0    3d22h
Total number of connected peers: 1
Total number of connections:     1
```

Example 1-54 *R8 DLSw+ Connectivity and* **show output**

```
dlsw local-peer peer-id 10.8.8.8 promiscuous
dlsw bridge-group 1
!
interface FastEthernet0/0
bridge-group 1
!
bridge 1 protocol ieee

R8#sh dlsw peer
Peers:                state     pkts_rx   pkts_tx  type  drops ckts TCP   uptime
 TCP 10.50.50.1      CONNECT      11346     11916  prom      0    0   0    3d23h
Total number of connected peers: 1
Total number of connections:     1
```

- *Set up a one-line filter, which only allows common SNA traffic to egress from R5 VLAN4 into the DLSw+ network.*

You are required to configure **lsap-output-list**, which allows only the SNA common traffic listed under access-list 200 to egress the DLSw+ connection to R8 on R5. If you have configured this correctly as in Example 1-55, you have scored 2 points.

Example 1-55 *R5 LSAP Filter Configuration*

```
dlsw remote-peer 0 tcp 10.8.8.8 lsap-output-list 200
!
access-list 200 permit 0x0000 0x0D0D
```

Section 7: IOS and IP Features (10 Points)

- *R2 is sited in a shared data center; make the serial link back into R1 as secure as possible at Layer 2.*

Did you think about IP Security (IPsec)? Well that is Layer 3; as is any form of access-list, you should realize that PPP is Layer 2 and this protocol has the capability to run CHAP over it, which makes the serial link very secure. PPP with CHAP over a serial link is just as happy as PPP over ISDN; the configuration is exactly the same. If you have configured this correctly as in Example 1-56 and Example 1-57, you have scored 2 points. Example 1-58 shows CHAP in action over the serial link.

Example 1-56 *R1 Serial Line PPP CHAP Configuration*

```
username disco password 0 cisco
!
interface Serial0/0
 ip address 10.90.90.2 255.255.255.240
 encapsulation ppp
 clockrate 2000000
 ppp authentication chap
 ppp chap hostname misco
```

Example 1-57 *R2 Serial Line PPP CHAP Configuration*

```
username misco password 0 cisco
!
interface Serial0/0
 ip address 10.90.90.1 255.255.255.240
 encapsulation ppp
 ppp authentication chap
 ppp chap hostname disco
```

Example 1-58 *R1* **debug ppp authentication** *Output*

```
R1#debug ppp authentication
PPP authentication debugging is on
3d23h: Se0/0 PPP: Treating connection as a dedicated line
3d23h: %LINK-3-UPDOWN: Interface Serial0/0, changed state to up
3d23h: Se0/0 CHAP: Using alternate hostname misco
3d23h: Se0/0 CHAP: O CHALLENGE id 2 len 26 from "misco"
```

continues

Example 1-58 *R1* **debug ppp authentication** *Output (Continued)*

```
3d23h: Se0/0 CHAP: I CHALLENGE id 5 len 26 from "disco"
3d23h: Se0/0 CHAP: Using alternate hostname misco
3d23h: Se0/0 CHAP: O RESPONSE id 5 len 26 from "misco"
3d23h: Se0/0 CHAP: I RESPONSE id 2 len 26 from "disco"
3d23h: Se0/0 CHAP: O SUCCESS id 2 len 4
3d23h: Se0/0 CHAP: I SUCCESS id 5 len 4
3d23h: %LINEPROTO-5-UPDOWN: Line protocol on Interface Serial0/0, changed state to
up
```

- *Ensure traffic from VLAN4, including any attached router interfaces to VLAN4, is hidden behind R5 Lo0 address when directed toward all external router networks.*

The question requires NAT configuration from VLAN4 R5 FastEthernet0/0 to all external networks egressing from R5 ATM3/0. The configuration is not complex so this should start the alarm bells ringing again. Where else did you hear them? The DLSw+ was not difficult so it may be worth checking as this does not break any previous connectivity. In fact, you should get into the habit of frequently checking your work to ensure your points are still safe.

With NAT enabled, R5's VLAN4 interface will NAT from the original source address 10.50.50.1 into source address 10.5.5.5; this will affect the DLSw+ peering that has been configured. When DLSW+ initiates, two connections are set up between the routers; after a capabilities exchange, the router with the lowest peer IP address has its connection dropped with the other maintained. Originally R5 (10.50.50.1) had a higher peer IP address than R8 (10.8.8.8), so it dropped the connection to R8. Now R5 has been NAT'd into source address 10.5.5.5; however, R5 still uses the source address 10.50.50.1 for its peering and drops the connection to R8 believing it is lower than its own source address. R8, as it is promiscuous, sees a connection coming in from 10.5.5.5 and then drops the connection, which is lower than its own peer address of 10.8.8.8. You can imagine that the DLSW+ will never peer correctly now with both peers dropping their connections as shown in Example 1-59. The only way to overcome this is by excluding the actual DLSw+ process from the NAT. To do this you will have to know that DLSw+ over TCP uses ports 2065 and 2067 for read and write; these should, therefore, not be NAT'd. You could gain the port information from an appropriate debug but ideally you should know this. If you have configured this correctly as in Example 1-60, you have scored 3 points.

Example 1-59 *R5* **debug dlsw** *Output*

```
R5#debug dlsw
4w0d: DLSw: START-TPFSM (peer 10.8.8.8(2065)): event:DLX-KEEPALIVE REQ
state:CONNECT
4w0d: DLSw: dtp_action_q() keepalive request from peer 10.8.8.8(2065)
4w0d: DLSw: Keepalive Response sent to peer 10.8.8.8(2065))
4w0d: DLSw: END-TPFSM (peer 10.8.8.8(2065)): state:CONNECT->CONNECT
4w0d: DLSw: dlsw_tcpd_fini() for peer 10.8.8.8(2065)
4w0d: DLSw: tcp fini closing connection for peer 10.8.8.8(2065)
```

Example 1-59 *R5* **debug dlsw** *Output (Continued)*

```
4w0d: DLSw: START-TPFSM (peer 10.8.8.8(2065)): event:ADMIN-CLOSE CONNECTION
  state:CONNECT
4w0d: DLSw: dtp_action_b() close connection for peer 10.8.8.8(2065)
4w0d: DLSw: END-TPFSM (peer 10.8.8.8(2065)): state:CONNECT->DISCONN
```

Example 1-60 *R5 Required NAT Modification Configuration*

```
interface FastEthernet0/0
 ip address 10.50.50.1 255.255.255.248
 ip nat inside
 interface ATM3/0
 ip address 10.99.99.2 255.255.255.248
 ip nat outside
!
ip nat inside source list 100 interface Loopback0 overload
!
access-list 100 deny   tcp host 10.50.50.1 eq 2065 host 10.8.8.8
access-list 100 deny   tcp host 10.50.50.1 eq 2067 host 10.8.8.8
access-list 100 deny   tcp host 10.50.50.1 host 10.8.8.8 eq 2065
access-list 100 deny   tcp host 10.50.50.1 host 10.8.8.8 eq 2067
access-list 100 permit ip 10.50.50.0 0.0.0.255 any
```

- *A router is to be installed onto VLAN4 in the future. This router will have a default configuration, so allow R6 to assist dynamically to aid the configuration process. The router will require an IP address of 10.50.50.6 and should load a configuration file called R9-config from a fictitious TFTP server on 172.16.0.59.*

The question requires that AutoInstall is used. This is a recent feature, which allows you to place a router with a default configuration onto a network, and it can be configured dynamically by receiving a DHCP address and TFTP server location for its own valid configuration. R6 is required to issue the DHCP address, TFTP server details, and **default-router** of R5 (10.50.50.1) that the router requires to contact the TFTP server on 172.16.0.59. The DHCP pool configuration on R6 excludes the majority of host addresses for network 10.50.50.0/29; this ensures that the only address offered to DHCP request is 10.50.50.6. You should notice that R6 is not connected to VLAN4 and AutoInstall works by DHCP request; for this reason, you must configure a **helper-address** on R5 to forward the request to R6. If you have configured this correctly as shown in Example 1-61 and Example 1-62, you have scored 5 points.

Example 1-61 *R6 Required Configuration for AutoInstall*

```
service dhcp
!
ip dhcp excluded-address 10.50.50.1 10.50.50.5
!
ip dhcp pool 1
   network 10.50.50.0 255.255.255.248
   bootfile R9-config
   option 150 ip 172.16.0.59
   default-router 10.50.50.1
!
```

Example 1-62 *R5 VLAN4 DHCP Relay Configuration*

```
interface FastEthernet0/0
 ip address 10.50.50.1 255.255.255.248
 ip helper-address 10.6.6.6
```

Section 8: QoS (8 Points)

- *Achieve maximum quality of voice calls by ensuring the real-time packet interval of 10 ms is not exceeded. Do not use RSVP in your solution.*

This question requires that fragmentation is used over the Frame Relay and ISDN networks that will transport voice traffic. You should remember that voice is still required should the Frame Relay network fail. By fragmenting the data, you can tailor the packet interval and ensure that voice quality is not compromised over low bandwidth links. Some basic math is required to calculate the current real-time packet interval over the Frame Relay and ISDN network to begin as detailed in Table 1-3. Note the Frame Relay speed is 256 kbps and the ISDN is 64 kbps using only one B channel.

Table 1-3 *MTU Values According to Bandwidth*

	Frame Relay (256 kbps)	ISDN (64 kbps)
No. Bytes TX'd per second	32,000	8,000
No. of bytes TX'd in 10ms	**320**	**80**

As Table 1-3 shows, 32,000 are bytes transmitted every second over the Frame Relay circuit (256,000 divided by 8) if the real-time delay or serialization delay is to be 10 ms; 320 bytes can be transmitted in this period (32,000 * 10 ms). Similarly, 80 bytes can be transmitted for the ISDN with the circuit speed of 64 kbps.

After you have calculated that 320 bytes will be transmitted over the Frame Relay network and 80 bytes over the ISDN in the 10 ms interval, you can adjust the interface MTU to reflect this for the ISDN and then change the **frame-relay fragmentation** to **320** under the VoFR map-class on both R1 and R4. If you have configured this correctly as shown in Example 1-63 and Example 1-64, you have scored 4 points.

Example 1-63 *R1 QOS MTU Configuration*

```
interface BRI0/0
 ip mtu 80
!
map-class frame-relay ccie
 frame-relay fragment 320
```

Example 1-64 *R4 QOS MTU Configuration*

```
interface BRI0/0
 ip mtu 80
!
map-class frame-relay ccie
 frame-relay fragment 320
```

- *To reduce the packet fragmentation in your network, allow R5 to determine appropriate fragmentation requirements when TCP sessions are originated from it to any part of the network.*

R5 should be configured with the global command **ip tcp path-mtu-discovery**. If you have configured this correctly, you have scored 2 points.

Section 9: Multicast (4 Points)

- *Enable your network to allow hosts on VLAN4 to receive and send multicast traffic from and to VLAN2; only perform configuration on R5 and R6 using PIM sparse dense mode.*

This question simply requires basic multicast setup between R5 and R6 using PIM **sparse-dense-mode**. If you have configured this correctly as shown in Example 1-65 and Example 1-66, you have scored 1 point.

Example 1-65 *R5 Multicast Configuration*

```
ip multicast-routing
!
interface FastEthernet0/0
 ip pim sparse-dense-mode
!
interface ATM3/0
 ip pim sparse-dense-mode
```

Example 1-66 *R6 Multicast Configuration*

```
ip multicast-routing
!
interface FastEthernet0/0
 ip pim sparse-dense-mode
!
interface ATM1/0.99 point-to-point
 ip pim sparse-dense-mode
```

- *Configure R6 to respond to pings from R5 to the multicast address of 224.4.4.4.*

Configure **ip igmp join-group 224.4.4.4** under R6 fastEthernet0/0. If you have configured this correctly, you have scored 1 point.

- *Do not allow R5 to fully participate in the PIM process by not allowing it to become a neighbor but do allow any IGMP messages generated by hosts on VLAN4 to be received by R6.*

This question requires Stub Multicast Routing. This allows you to configure remote/stub routers as IGMP proxy agents. The stub router does not fully participating in PIM and, hence, is not seen as a PIM neighbor; access-list 11 blocks the neighbor. If you have configured this correctly as in Example 1-67, you have scored 2 points.

Example 1-67 *R6 Stub Multicast Configuration*

```
interface ATM1/0.99 point-to-point
 ip pim neighbor-filter 11
!
access-list 11 deny   10.99.99.2
```

How Did You Do?

With the aid of the answers section, full configurations, and routing tables on the CD, you should now have an accurate evaluation of your lab. If you scored more than 80 points within the time frame, you should congratulate yourself; you are well on the way to becoming a CCIE, you have demonstrated the ability to think laterally and shown an impressive knowledge of your subject. If you scored less than 80 don't worry, this will improve as you progress throughout the book and learn how to analyze each question and improve your core skills.

Did you spot the landmines in the lab? The classics in Lab 1 where EIGRP neighbor issues, the ISDN line staying up, the BGP requiring policy routing, and the NAT breaking DLSw+, ideally you should be able to spot these before configuration and factor them in; if not, it shows you how important it is to read the paper thoroughly and ensure everything works over and over again after configuration.

This is not to say that in the real exam there will be any items that could catch you out, but by being on your guard, you will ensure that your quality of work is far higher.

You might feel that the questions were too vague or you did not have sufficient time to complete the lab but this is what you will be met with when you open your folder containing your real exam at the test center; your ability to spot landmines (if present), ask the right questions, and configuration speed will improve as you tackle each practice lab.

For each question that you did not answer correctly, take the time to research the subject thoroughly and turn your weaknesses into your strengths. This with plenty of practice is how ultimately you will gain your number.

Further Reading/URLs

To better prepare yourself and follow up on the topics covered in this lab, you should consult the http://www.cisco.com/public/pubsearch.html website for more information on the following topics:

- 3550 802.1X
- EIGRP Stub Routing
- Dialer-Watch
- PPP Link Quality
- DLSW Filtering SNA
- DLSw and NAT
- AutoInstall
- PIM Neighbor Filtering