CHAPTER 1

Practice Lab 1

All labs in this book are multi-protocol, multi-technology, testing you in areas such as Routing, Switching, Security, and VPN, as outlined in the CCIE Security blueprint. When you first read the questions in the lab, you might find them fairly easy, but they are carefully written to present high complexity and many hidden problems. Such is the case in the real CCIE lab exam.

To assist you, solutions are provided for the entire lab, including configurations and common **show** command outputs from all the devices in the topology. Furthermore, a "Verification, Hints, and Troubleshooting Tips" section is provided, which gives you tips and hints to troubleshoot and identify the hidden problem or trick in the question.

This is the first lab of seven in this book. Each lab is 8 hours and weighs 100 marks, passing of which is 80 marks. The objective is to complete the lab within 8 hours and obtain a minimum of 80 marks to pass. This test has been written such that you should be able to complete all questions, including initial configuration (such as IP addressing), within 8 hours; this excludes cabling time. Allow up to 1 hour for cabling, use the cabling instructions, and observe the instructions in the general guidelines. You can use any combination of routers as long as you fulfill the topology diagram in Figure 1-1. It is not compulsory to use the same model of routers.

NOTE Cabling and IP addressing are already completed on the real CCIE Lab. You are not required to do any cabling or the IP addressing.

Equipment List

- 6 routers with the following specifications (all routers are to be loaded with the latest Cisco IOS version in 12.1(T) train):

  R1 — 4 serial, 1 BRI (with IP Plus image)

  R2 — 2 serial, 1 Ethernet (with IP Plus + Firewall image)

  R3 — 2 serial, 1 Ethernet, 1 BRI (with IP Plus + IPSec 56 image)

  R4 — 1 serial, 1 Ethernet (with IP Plus + Firewall + IPSec 56 image)

  R5 — 1 serial, 1 Ethernet (with IP Plus image)

  R6 — 5 serial, 3 Ethernet (with IP Plus + IPSec 56 image)

- 1 switch 3550
- 1 PIX — 2 interfaces (with version 6.x)
- 1 PC with Windows 2000 Server with CiscoSecure ACS 3.x+
- The IDS device in the topology is not required; it is there to give you an idea to configure other aspects of this lab. Subsequent chapters do require a Network IDS appliance.

Figure 1-1 *Lab Topology Diagram*

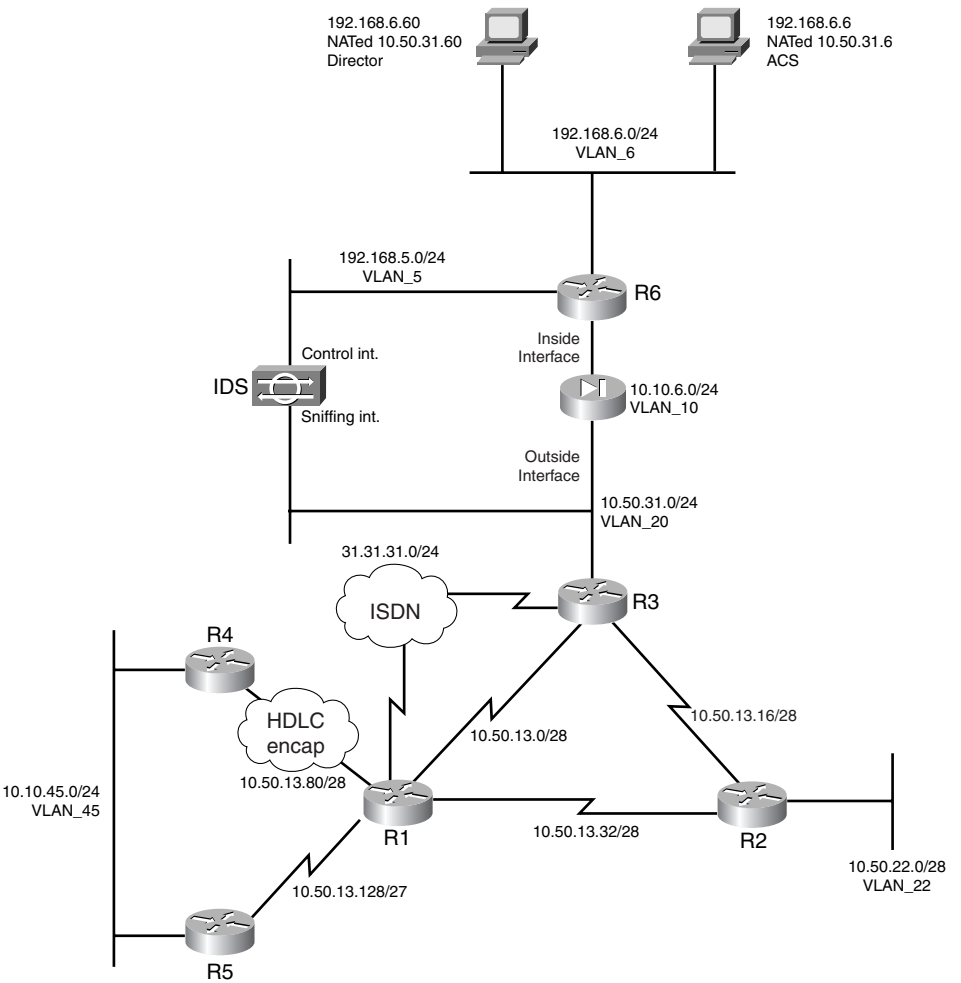

General Guidelines

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

- Use DLCIs provided in the diagram.
- Use the IP addressing scheme provided in the diagram; do not change any IP addressing unless otherwise specified. In the CCIE Lab, initial configurations are loaded, and therefore IP addresses are not to be changed. In this book, each chapter has a separate lab topology with different IP addressing, so each chapter needs to be recabled and all IP addresses need to be redone from the previous chapter.
- Use **cisco** as the password for any authentication string, enable-password, and TACACS+/RADIUS key or for any other purpose.
- Add additional loopbacks as specified during this lab.
- Configure VLANs on Switch1 as per Figure 1-1.
- All routers should be able to ping any interface in the network using the *optimal* path.
- You must time yourself to complete this lab in 8 hours.
- Do not use any external resources or answers provided in this book when attempting the lab.
- Configure a backdoor for any of the AAA questions below to the local database. If you don't, you will lose all points for that question.
- Do not configure any authentication or authorization on the console and aux ports.

Setting Up the Lab

You can use any combination of routers as long as you fulfill the topology diagram outlined in Figure 1-1. It is not compulsory to use the same model of routers.

Frame Relay DLCI Information

Only DLCIs indicated in Figure 1-2 should be mapped on the routers.

Figure 1-2 *Frame Relay DLCI Diagram*

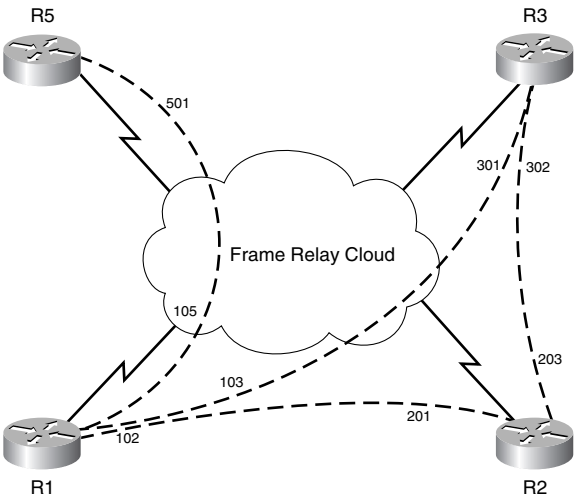

Routing Protocol Information

Use Figure 1-3 to configure routing protocols for the exercises to follow.

Figure 1-3 *Routing Protocol Information*

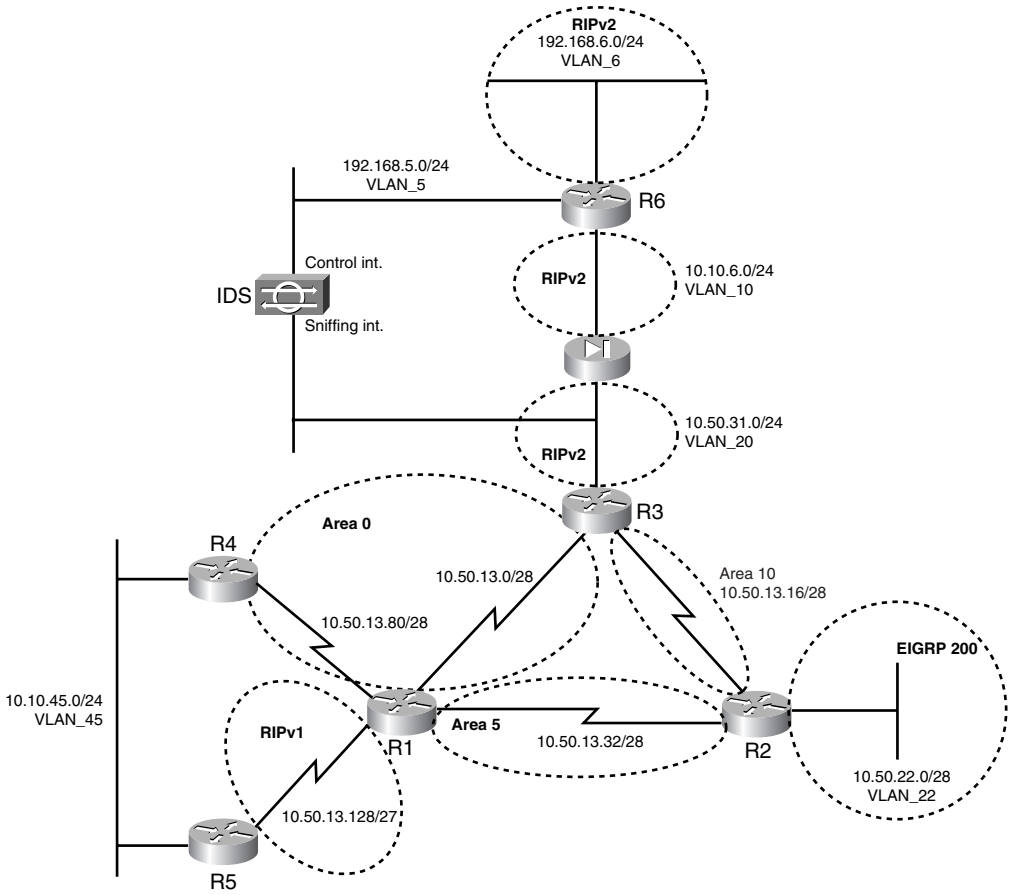

BGP Information

Use Figure 1-4 to configure BGP.

Cabling Instructions

Use Tables 1-1 and 1-2 for cabling all devices in your topology. It is not a must to use same type or sequence of interface. You may use any combination of interface(s) as long as you fulfill the requirement.

Figure 1-4 *BGP Information*

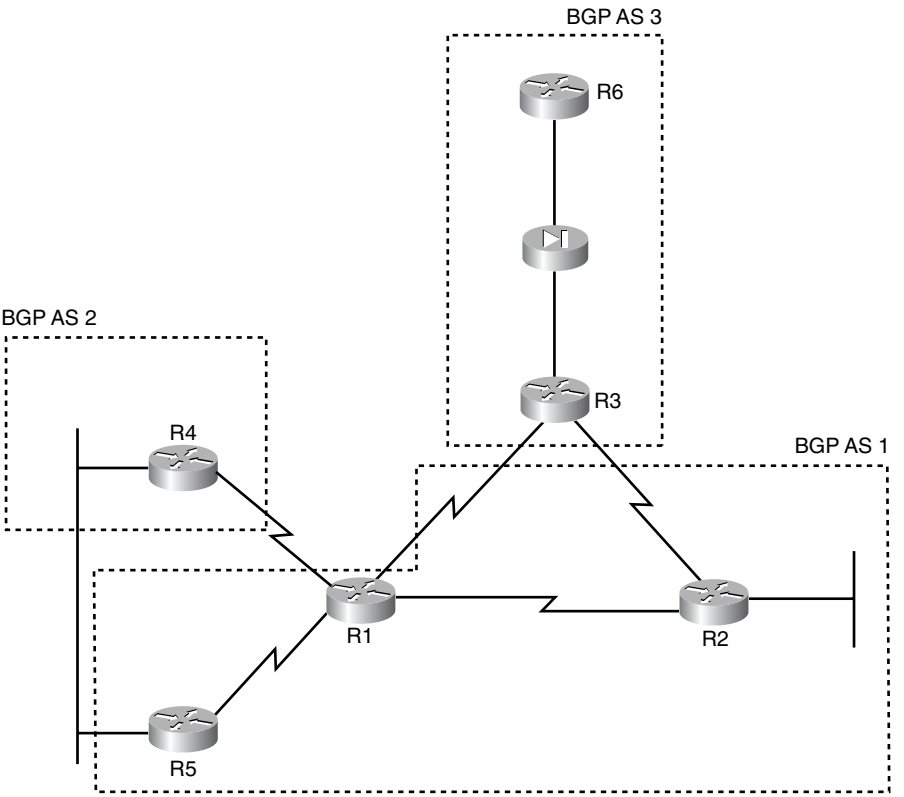

Table 1-1 *Cabling Instructions (Ethernet)*

Ethernet Cabling	Switch1
R2-ethernet0/0	Port 1
R3-fastethernet0/0	Port 2
R4-fastethernet2/0	Port 3
R5-ethernet0	Port 4
R6-ethernet0/0 to VLAN6	Port 5
R6-ethernet0/1 to PIX	Port 6
PIX-inside-ethernet1	Port 7
PIX-outside-ethernet0	Port 8

continues

Table 1-1 *Cabling Instructions (Ethernet) (Continued)*

Ethernet Cabling	Switch1
IDS-control-interface	Port 9
IDS-sniffing-interface	Port 10
AAA server	Port 11

Table 1-2 *Cabling Instructions (Serial)*

Back-to-Back Cabling	DTE-End	DCE-End
R1-to-frsw* (to R2/R3)	R1-serial2/0	R6-serial1/0
R1-to-frsw* (to R5)	R1-serial2/1	R6-serial1/1
R1-to-R4	R1-serial2/2	R4-serial3/0
R2-to-frsw* (to R1/R3)	R2-serial1/0	R6-serial1/2
R3-to-frsw* (to R1/R2)	R3-serial1/0	R6-serial1/3
R5-to-frsw* (to R1)	R5-serial0	R6-serial1/4

*frsw = Frame Relay Switch R6

Practice Lab 1 Exercises

Section 1.0: Basic Configuration (10 points)

1.1: IP Addressing (2 points)

1. Redraw a detailed topology with all necessary information.
2. Configure IP addressing as per the diagram.
3. Do not configure any static or default routes anywhere on the network unless otherwise specified. Configure a default route on R2. Your routing table should have an entry as follows: “Gateway of last resort is 0.0.0.0.” Populate this default route to all the routers.
4. Create the following loopbacks:

   Loopback-1 11.11.11.11/24 on R1

   Loopback-2 111.111.111.111/24 on R1

   Loopback-1 12.12.12.12/24 on R2

   Loopback-2 122.122.122.122/24 on R2

   Loopback-1 13.13.13.13/24 on R3

Loopback-2 133.133.133.133/24 on R3

Loopback-3 192.168.3.1/24 on R3

Loopback-1 14.14.14.14/24 on R4

Loopback-2 144.144.144.144/24 on R4

Loopback-1 16.16.16.16/24 on R6

Loopback-2 166.166.166.166/24 on R6

1.2: Frame Relay Configuration (4 points)

1. Configure R6 as a Frame Relay switch. Use the DLCI information provided for Frame Relay routing as per Figure 1-2.
2. Configure Frame Relay between R1, R2, R3, and R5. Configure point-to-point subinterfaces on all routers. Do not configure a subinterface on R1 for serial connection to R5. Use only the DLCIs provided in the DLCI information diagram. Use LMI type Cisco. The speed should be set to 56 KB on the DCE ends.

1.3: LAN Switch Configuration (4 points)

1. Configure Switch1 with the VLAN information provided in the diagram shown in Figure 1-1. Also make sure that it is easier for the network administrator to troubleshoot port/VLAN identification.
2. Configure security such that network devices are operational on allocated ports only. In the event of a security breach, the administrator should take strict action.
3. Configure the management interface of the switch with IP address 10.10.45.45. Only R4, R5, and R1 should have Telnet access to the switch. Configure redundancy such that the management interface is reachable from R1 if the serial link is down between R1 and R4. Configure a static route on R1 for the 10.10.45.0/24 network. Do not configure any routing protocol on Switch1 to achieve this task; you can use static routes as required.
4. Configure port 9 on the switch to be in VLAN 5. There is an IDS sensor deployed off r6. It has been preconfigured. The aim is to protect the PIX outside interface, so configure accordingly.

Section 2.0: Routing Configuration (25 points)

2.1: Core Routing OSPF/EIGRP/RIP (5 points)

1. Configure OSPF, EIGRP, and RIP as shown in Figure 1-3. All routing/update traffic should be encrypted. Mutually redistribute between IGPs only where necessary.

2.2: OSPF (4 points)

1. Configure a loopback on R3 10.50.13.97/28 in Area 66. R5 should see this network in the routing table. Do not use any summarization technique to achieve this task. Performance should not be compromised.
2. Configure the following loopbacks on R3; put them in Area 30 on R3.

   30.30.1.0/24

   30.30.2.0/24

   30.30.3.0/24

   30.30.4.0/24

   30.30.5.0/24

   30.30.6.0/24

2.3: EIGRP (3 points)

1. Configure three null routes on R2 to appear in the EIGRP-200 database for the following subnets: 10.50.22.16/28 10.50.22.32/28 10.50.22.64/28. Redistribute EIGRP-200 on R2 into OSPF. All other routers should see these routes as *one* route with a cost of 10.

2.4: RIP (3 points)

1. Configure RIPv2 on PIX to peer with inside router R6 and outside router R3. Use strong encryption. Do not configure a static default route. PIX should learn all routes via RIP. You must ensure that no other device can establish adjacency with the PIX and that routing updates are secured.
2. Configure RIP version 1 between R1 and R5. R5 should be able to ping all parts of the network.
3. Advertise VLAN 6 network 192.168.6.0/24 on R6 in RIPv2. Make sure you can ping the AAA server from the PIX.

2.5: BGP (10 points)

2.5.1: Basic BGP Configuration (2 points)

1. Configure the BGP peers as follows using Figure 1-4.

   R2 – R3 eBGP

   R2 – R1 iBGP

   R1 – R3 eBGP

   R1 – R5 iBGP

R5 – R4 eBGP

R3 – R6 iBGP (configure static NAT 10.10.6.2 to 10.50.31.22 on PIX to achieve this task)

NOTE You can use "no sync" on all BGP peers.

2.5.2: BGP Connections (2 points)

1 Ensure that eBGP connection state on R6 shows local port as 179 always.

2.5.3: BGP and OSPF (2 points)

1 Advertise loopback2 on R2 and R4 in BGP. Redistribute BGP into OSPF on these routers so that BGP routes on all OSPF routers are seen as OSPF (E1) and not through BGP. Ensure all routers can ping these loopbacks using the optimal path. Do not use the **distance** command to achieve this task.

2.5.4: BGP and RIP (2 points)

1 Advertise loopback2 on R1 in BGP and RIPv1. Advertise loopback1 in RIPv1 only. R5 should be able to ping all routers in the network and vice versa.

2.5.5: BGP Attributes (2 points)

1 Advertise loopback1 and loopback2 on R6 in BGP. Do not use the **network** statement to advertise loopback2. R3 should see both loopbacks as internal. Ensure all routers in the network can ping these loopbacks using the optimal path.

Section 3.0: ISDN Configuration (8 points)

3.1: Basic ISDN (4 points)

1 Configure ISDN on R1 and R3. Use network 31.31.31.0/24. Advertise this network in OSPF area0.

ISDN information:

BRI number on R3 is 99281766

BRI number on R1 is 99281764

Switch-type = basic-ts013

2 Configure redundancy such that if there is a change in backbone database on R3, BRI comes up and all traffic continues normally.

3.2: PPP Callback (4 points)

1 R1 should call back R3 using TACACS+. Configure static NAT translation on PIX for AAA server behind R6 to achieve this task, as shown in topology diagram Figure 1-1.

Section 4.0: PIX Configuration (5 points)

4.1: Basic PIX Configuration (2 points)

1 Configure PIX inside and outside interface 10.10.6.1 and 10.50.31.1 respectively. Do not configure a default route on PIX. All routes should be learned via RIP as per Section 2.4.

2 You should be able to ping all routers in the network from the PIX, including the AAA server and R6 networks behind the PIX.

4.2: Network Address Translation (NAT) (2 points)

1 Configure static NAT translation and an access list on PIX to receive reliable syslog messages for a server behind R6. NAT 192.168.6.65 as 10.50.31.65.

4.3: Advanced Configuration (1 point)

1 A workstation on VLAN 6 is failing to ping a server on the same VLAN. Both PCs are in the same VLAN. Upon investigating, it is determined that the workstation is seeing the PIX inside MAC address for the server. When ethernet1 on PIX is shut down, the workstation can ping successfully. Resolve this without shutting down the ethernet1.

Section 5.0: IPSec/GRE Configuration (10 points)

5.1: IPSec (5 points)

5.1.1: IPSec LAN-to-LAN Using Preshared (2 points)

1 Encrypt IDS traffic between PIX and R4 in Section 6.2.1. Use a preshared key and SHA for message authentication and DES for encryption. Configure all other parameters as you feel appropriate.

5.1.2: Advanced IPSec LAN-to-LAN (3 points)

1. Configure IPSec to encrypt GRE traffic between R6 and R3 in Section 5.2.
2. Use preshared keys. Configure all other parameters as you feel appropriate.
3. If there is a loss of connectivity between two IPSec peers, terminate the sessions.
4. You are allowed to put one static route but not a default route on the PIX to achieve this task.

5.2: GRE (5 points)

1. Configure GRE through PIX; R6 should see all loopbacks in area 30 created on R3 in Section 2.2. R6 should ping even networks through GRE and odd networks through PIX. Run EIGRP-100 on GRE. Any ACL used to accomplish this task should not be more than one line.

Section 6.0: IOS Firewall + IOS IDS Configuration (10 points)

6.1: CBAC (6 points)

6.1.1: Basic CBAC Configuration (2 points)

1. Configure IOS Firewall on R2 to protect the EIGRP network. Ensure it can reach the rest of the network.

6.1.2: Firewall Filtering (2 points)

1. No access but ICMP is allowed to R2.
2. R1 should be able to Telnet to R2 using its loopback2 address as source. Configure ingress ACL on WAN links, including anti-spoofing technique. Do not deny RFC1918 address space.

6.1.3: Advanced CBAC Configuration (2 points)

1. Configure prevention against TCP host-specific denial-of-service on R2. Set the threshold to 200 before the firewall engine starts deleting half-open sessions to the host.

6.2: Intrusion Detection System (IDS) (4 points)

6.2.1: Basic IDS Configuration (2 points)

1. Configure IDS on R4 to protect the Ethernet network from internal intrusion, and configure to send an alarm for info and attack matching signatures.
2. Use the following details:

   Director Host id 5, Sensor Host-id 4

   Org id 100, Org name cisco

   Director IP is 192.168.6.60 (create NAT on PIX to 10.50.31.60 to achieve this task)

6.2.2: Signature Tuning (1 point)

1. The message in the following line is received on the syslog server:

```
Jun 28 10:52:25.538: %IDS-4-TCP_SYN_ATTACK_SIG: Sig:3050:Half-Open Syn
   Flood - from 10.50.16.5 to 144.144.144.144
```

   Upon investigation it was discovered that there is a specific application running on this machine. Consider these as false alarms; configure the IDS not to send such alarms in the future.

6.2.3: Spam Attack (1 point)

1. R4 is experiencing a spam attack. An alarm should only be generated if the spam attack has more than 500 recipients in a mail message.

Section 7.0: AAA (7 points)

7.1: AAA on the Router (4 points)

1. Configure router authentication and authorization on R4 using TACACS+. Configure two users on ACS, "user1" and "user2." User1 should have privilege level 10 and user2 privilege level 15. Configure such that User1 is able to run the command **show running-configuration** only, and user2 is able to run all commands.
2. Configure redundancy such that in the event the TACACS+ server is down, both users are able to log in using the local database and maintaining the same authorization.
3. When user1 or user2 logs in, they should get the # prompt with their respective privilege level without entering the **enable** command.

4 Configure fallback to local in the event the AAA server is down. Do not configure any authentication or authorization for console and auxiliary ports.

7.2: AAA on PIX (3 points)

1 Users should be able to Telnet to R6 loopback1 from anywhere on the network. Configure username "r6telnet" on ACS with the necessary parameters. Configure authentication and port authorization on PIX to achieve this task.

Section 8.0: Advanced Security (10 points)

8.1: Password Protection (2 points)

1 Make sure when users see the configuration of the router, all passwords are secured and not readable.

8.2: EXEC Authentication (4 points)

1 Encrypt the enable password on R2 with a nonreversible algorithm denoted by the number 5 in the configuration.

2 R2 should prompt for a username/password for privilege access and authenticate with the TACACS server. Do not use any AAA commands to achieve this task. In the event when the TACACS server is down, allow users to log in successfully. Do not use the **tacacs-server last-resort** command to achieve this task.

8.3: Access Control (4 points)

1 Configure such that a username **testconfig** with password **testconfig** is able to see the current configuration of R3 from anywhere on the network without having login access to the router.

2 Configure R5 vty line so that only loopback2 of R3 is able to Telnet.

Section 9.0: IP Services and Protocol-Independent Features (10 points)

9.1: NAT (4 points)

1 Create a loopback on R3 with 192.168.3.1/24. Configure NAT translation on R3 for this network to be translated to interface IP address with overload. You should be able to

ping anywhere in the network from R3 sourcing from this loopback and get NATed to the corresponding egress interface.

9.2: NTP (2 points)

1 Configure R1 clock polling from NTP server R2. All NTP packets should be encrypted. Update the system.

9.3: SNMP (2 points)

1 Configure R3 to report the BGP configuration to Network Management System 192.168.6.99 (NATed 10.50.31.99). Configure the appropriate static/ACL on the PIX to achieve this task.

9.4: Policy Routing (2 points)

1 There is a mail server 10.50.31.98 and a web server 10.50.31.99 on VLAN20. Configure such that networks behind and from R1

Traverse via R2 to reach the mail server

Traverse via R3 to reach the web server

Section 10.0: Security Violations (5 points)

10.1: Denial of Service—DoS (3 points)

1 R3 is experiencing an ICMP DoS attack on the WAN links. Take necessary action to prevent this. Do not deny ICMP.

10.2: IP Spoofing (2 points)

1 Configure PIX to perform a route lookup based on the source address to protect from an IP spoofing attack using network ingress and egress filtering, as described in RFC 2267.

Verification, Hints, and Troubleshooting Tips

As mentioned in the Overview, this section is primarily important when you're configuring the exercise and it is not working for you. You can use this section to verify and compare results by adapting the troubleshooting methods shown. Also provided are hints needed to configure and complete the respective exercises. Sometimes, it is easy to misinterpret the question, which has hidden and tricky elements required to be configured.

This section also guides you in using the most common **show** and **debug** commands used for verification and troubleshooting, which are very handy. Learn to use them and to read the outputs and when to use them.

There are several key elements to pass the exam, one of which is to troubleshoot effectively.

Section 1.0: Basic Configuration

1.1: IP Addressing

1. Configure IP addresses as per the topology diagram shown in Figure 1-1.
2. Configure all the loopbacks and advertise them as per the instructions in different sections of the exercise.
3. Configure a default route on R2 to Ethernet 0/0. This will show "Gateway of last resort is 0.0.0.0" in your routing table. Propagate the default route to all other routers.
4. Configure **default-information originate always** on R2. Do not configure any static routes unless otherwise specified.

1.2: Frame Relay Configuration

1. Map only DLCIs specified in the diagram. Do **show frame-relay map** and check to see if there are any additional DLCIs dynamically populated that are not required. If so, turn off inverse-arp on that interface using the **no frame-relay inverse-arp** command.

1.3: LAN Switch Configuration

1. Configure IP address 10.10.45.45 on VLAN 45.
2. Configure a port description on the interface for identification.
3. Configure VLAN names as per the topology diagram.
4. Configure port security on all ports except port 10 (span destination port).

5 Configure a default route to 10.10.45.4 and a floating static route to 10.10.45.5 with higher admin distance for redundancy.

6 Configure an access list to permit R4, R5, and R1 and apply to vty lines. Note that you have to put two host entries for R1 for redundancy, one through R4 and another through R5; see switch configuration in the Solutions section. Test by sourcing the Telnet with Serial 2/1 and Serial 2/2 from R1 as follows:

```
r1#telnet 10.10.45.45 /source-interface Serial 2/1
r1#telnet 10.10.45.45 /source-interface Serial 2/2
```

7 Configured SPAN session and specify ports to monitor: source port 8 (PIX outside interface) and destination port 10 (sniffing interface).

8 Static route for 10.10.45.0/24 network on R1 should not be seen on any other routers.

Section 2.0: Routing Configuration

2.1: Core Routing OSPF/EIGRP/RIP

1 Configure core routing for all the above protocols on all routers in the network. Redistribute only where necessary; you must use your judgement. See the solutions below where required.

2.2: OSPF

1 There is a loopback on R3 10.50.13.97/28 in Area66. You will see this route on all participating OSPF routers but not on R5, which is running RIPv1. The RIP network between R5 and R1 is a /27. You are redistributing OSPF into RIP on R1, but it will not redistribute the /28 since it has a different mask belonging to the same major net.

2 The workaround is to summarize this in RIP on R1 to /27. Since you are restricted not to use a summarization technique, another technique to achieve this is to create a loopback on R1 with a /27 mask; this will automatically get into the RIP database, as it is the same major net. You can check this in the RIP database with **show ip rip database** on R1, and you will find the loopback as directly connected and not redistributed. This is not a good practice in real life but you can use this in lab setup.

3 It is always a good idea to hardcode the router IDs on each OSPF speaker. This way, it is easier to identify the router sending/receiving updates and troubleshoot any problems with the OSPF peers.

4 On R5, you need to enable **split-horizon** on the serial link to R1, as it is advertising the 30.0.0.0/8 route back to R1 (see the following debug output):

```
r1#debug ip routing
04:08:11: RT: network 144.144.0.0 is now variably masked
04:08:11: RT: add 144.144.0.0/16 via 10.50.13.130, rip metric [120/3]
04:08:11: RT: network 122.0.0.0 is now variably masked
04:08:11: RT: add 122.0.0.0/8 via 10.50.13.130, rip metric [120/3]
```

```
04:08:11: RT: network 13.0.0.0 is now variably masked
04:08:11: RT: add 13.0.0.0/8 via 10.50.13.130, rip metric [120/3]
04:08:11: RT: network 133.133.0.0 is now variably masked
04:08:11: RT: add 133.133.0.0/16 via 10.50.13.130, rip metric [120/3]
04:08:11: RT: network 30.0.0.0 is now variably masked
04:08:11: RT: add 30.0.0.0/8 via 10.50.13.130, rip metric [120/3]

04:08:39: RT: metric change to 144.144.0.0 via 10.50.13.130, rip metric [120/3]
        new metric [120/4]
04:08:39: RT: metric change to 122.0.0.0 via 10.50.13.130, rip metric [120/3]
        new metric [120/4]
04:08:39: RT: metric change to 13.0.0.0 via 10.50.13.130, rip metric [120/3]
        new metric [120/4]
04:08:39: RT: metric change to 133.133.0.0 via 10.50.13.130, rip metric [120/3]
        new metric [120/4]
04:08:39: RT: metric change to 30.0.0.0 via 10.50.13.130, rip metric [120/3]
        new metric [120/4]
```

Fix this by enabling split-horizon on Serial0 on R5.

Snip from R5 config:

```
interface Serial0
  ip address 10.50.13.130 255.255.255.224
  ip split-horizon
  no frame-relay inverse-arp
```

2.3: EIGRP

1. Configure EIGRP with the null routes, redistribute EIGRP into OSPF with a metric of 10, and then summarize them in OSPF to a /25 to advertise one route to all OSPF neighbors.

2.4: RIP

1. Basic RIP configuration to be done on R3, R6, and PIX using MD5 authentication. Do not configure any default route on PIX and R6. You should configure RIP on PIX to inject a default route for R3 using **rip inside default version 2 authentication md5 cisco 1**.
2. Make sure you can ping the AAA server from the PIX.

2.5: BGP

2.5.1: Basic BGP Configuration

1. Configure R1 as route-reflector server for BGP connection to R5, as it is not fully meshed. Also configure next-hop-self for R5 peer, as R1 will advertise all routes learned by iBGP peers and forward to R5 without changing the next hop, and this could cause reachability problems at times if you don't have proper routes on R5.

2. For iBGP between R3 and R6, you need to create static NAT for R6 Ethernet 10.10.6.2 to 10.50.31.22 and permit TCP port 49 on PIX for inbound connections. You will use 10.50.31.22 on R3 for BGP peer configuration.
3. It is always a good idea to hardcode the router IDs on each BGP speaker just like in the OSPF process for troubleshooting.

2.5.2: BGP Connections

1. This is a tricky one. The objective is to always build a BGP connection from outside-to-inside only. That is, R6 should not be able to build a BGP connection to R3, which it can by default since packets are going from a higher security level interface to a lower interface.
2. To achieve this task, you need to configure an ACL on the inside interface and deny R6 BGP connection to R3:

```
access-list inside deny tcp host 10.10.6.2 host 10.50.31.2 eq bgp (hitcnt=4)
```

   See also the PIX output in the Solutions section.

2.5.3: BGP and OSPF

1. Another tricky question. It seems very straightforward to create loopbacks on R2 and R4. Advertise them in BGP using the **network** command, and redistribute into OSPF. All OSPF routers should see these routes.
2. Well, it is not so simple. After doing the above, check from R3 and R4 and see if you can ping using the optimal path. If you do traceroute from routers R3 and R4, you will notice that they are not taking the optimal path.

   For example:

   Traceroute 122.122.122.122 from R4 and you will find that it is going through R5, whereas the optimal path is via R1.

   Traceroute 144.144.144.144 or 122.122.122.122 from R3 and you will find that it is not using the optimal path either.
3. The solution is to use the BGP **network backdoor** command. You can also use the **distance** command, but you are restricted to use this to achieve this task. See solutions for R3 and R4.

NOTE Note that there is a bug with the BGP **network backdoor** command in Cisco IOS Software Release 12.1T. When you apply the command it will not take effect. See bug id CSCdr12571 for more details. The workaround is to remove using the **no network backdoor** command and reapply back in. No need to clear BGP.

2.5.4: BGP and RIP

1. Traceroute 111.111.111.111 (loopback2) from R4. You will notice that the next hop is R5 10.10.45.5 and not R1 10.50.13.81 as it is for 11.11.11.11 (loopback1). Why?

Because we advertised 111.111.111.111 in RIP and BGP on R1, which made to BGP table on R5, and since R5 was peering eBGP with R4, it overwrote the route learned on R4 via OSPF as better admin distance. To confirm, turn on **debug ip routing** on R4 and **clear ip route *** as demonstrated in the following ip routing debug snippet from R4.

```
5d23h: RT: add 111.111.111.0/24 via 10.50.13.81, ospf metric [110/5]
5d23h: RT: closer admin distance for 111.111.111.0, flushing 1 routes
5d23h: RT: add 111.111.111.0/24 via 10.10.45.5, bgp metric [20/0]
```

As you can see, it overwrites the 111.111.111.111 route learned via OSPF with BGP.

To fix this, you need to tweak the eBGP admin distance from 20 to 120 (something higher than OSPF) as follows.

Configure the **distance bgp 120 200 200** command on R4 in BGP.

Then **clear ip route *** and you will see that the OSPF route stays, as demonstrated in the following ip routing debug from R4.

```
5d23h: RT: add 111.111.111.0/24 via 10.50.13.81, ospf metric [110/5]
```

NOTE You will not notice this problem if you have BGP "synchronization" enabled on R4; do "no sync" on R4 BGP and you will run into this problem.

2.5.5: BGP Attributes

1 Advertise loopback1 using the **network** command on R6. You are restricted to use the **network** command to advertise loopback2; you will need to redistribute connected in BGP. Create an access list and a route map to redistribute loopback2 only. After doing so, do a **show ip bgp** on R6 and you will find that the origin-code for loopback2 is incomplete, denoted by a **?**, because it has been redistributed and BGP hasn't learned this internally. To change the origin-code to denote **i**, use the **set origin igp** command in your route map:

```
access-list 16 permit 166.166.166.0 0.0.0.255
!
route-map loop2 permit 10
 match ip address 16
 set origin igp
!
router bgp 3
 no synchronization
 bgp router-id 6.6.6.6
 bgp cluster-id 2795939494
 bgp log-neighbor-changes
 network 16.16.16.0 mask 255.255.255.0
 redistribute connected metric 2 route-map loop2
```

Done? No, not yet.

Check the routing table on R3 to see if you see the R6 Loopback1 and Loopback2. Check its next hop, and which routing protocol is it learning from:

```
r3#show ip route
     16.0.0.0/24 is subnetted, 1 subnets
O E2    16.16.16.0 [110/5] via 10.50.13.1, 00:00:07, Serial1/0.1
     166.166.0.0/24 is subnetted, 1 subnets
O E2    166.166.166.0 [110/5] via 10.50.13.1, 00:00:07, Serial1/0.1
O*E2 0.0.0.0/0 [110/1] via 10.50.13.17, 00:00:07, Serial1/0.3
```

See the following debugs, which show Loopback1 and Loopback2 being learned via OSPF (10.50.13.1) and not iBGP (10.50.31.22). This is because R3 is peering eBGP with R1 and redistributing BGP into OSPF on R1. Because OSPF is peering with R3, it is learning the route via OSPF and overwriting the iBGP route learned via R6. Very complex loop!

```
r3#debug ip routing
IP routing debugging is on
r3#
4d23h: RT: closer admin distance for 16.16.16.0, flushing 1 routes
4d23h: RT: add 16.16.16.0/24 via 10.50.13.1, ospf metric [110/5]
4d23h: RT: add 122.122.122.0/24 via 10.50.13.17, ospf metric [110/786]
4d23h: RT: add 144.144.144.0/24 via 10.50.13.1, ospf metric [110/834]
4d23h: RT: closer admin distance for 166.166.166.0, flushing 1 routes
4d23h: RT: add 166.166.166.0/24 via 10.50.13.1, ospf metric [110/5]
```

To fix this, create an access list and filter the two loopbacks using distribute-list inbound in OSPF:

```
r3# show running-config
! <snip>
router ospf 110
 router-id 3.3.3.3
distribute-list 16 in
!
r3#show access-lists
Standard IP access list 16
    deny   16.16.16.0, wildcard bits 0.0.0.255
    deny   166.166.166.0, wildcard bits 0.0.0.255
    permit any

r3#
r3#debug ip routing
IP routing debugging is on
r3# <snip>
4d23h: RT: add 16.16.16.0/24 via 10.50.31.22, bgp metric [200/2]
4d23h: RT: add 166.166.166.0/24 via 10.50.31.22, bgp metric [200/2]
```

Routes are now being learned via iBGP and not OSPF.

Ping and traceroute to verify.

Section 3.0: ISDN Configuration

3.1: Basic ISDN

1 Configure legacy BRI on R1 and R3. Configure OSPF demand circuit for redundancy on R3.

3.2: PPP Callback

1 Configure R1 as callback server and R3 as callback client. Do not configure dialer-map on R1, as it will retrieve the callback number from the AAA server.

2 Configure AAA server with username "r3" and its callback attributes. Refer to Figure 1-5 for PPP callback user profile settings on ACS.

3 As a fallback, configure PPP authentication to local and a username "r3" on R1 with callback string:

```
aaa new-model
aaa authentication ppp default group tacacs+ local
aaa authorization network default group tacacs+ local
!
username r3 callback-dialstring 99281766 password 7 094F471A1A0A
```

Figure 1-5 *PPP Callback Settings on CiscoSecure ACS*

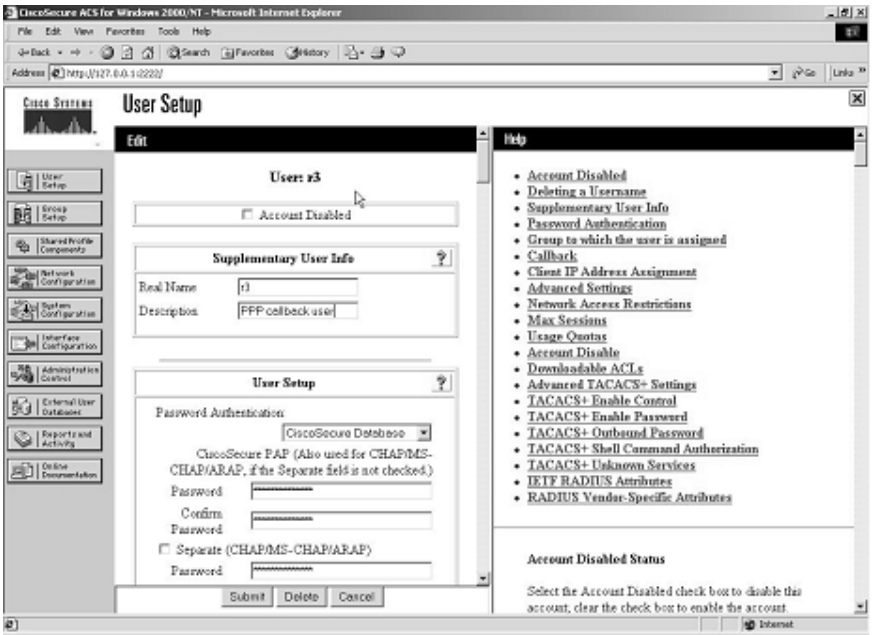

Figure 1-5 *PPP Callback Settings on CiscoSecure ACS (Continued)*

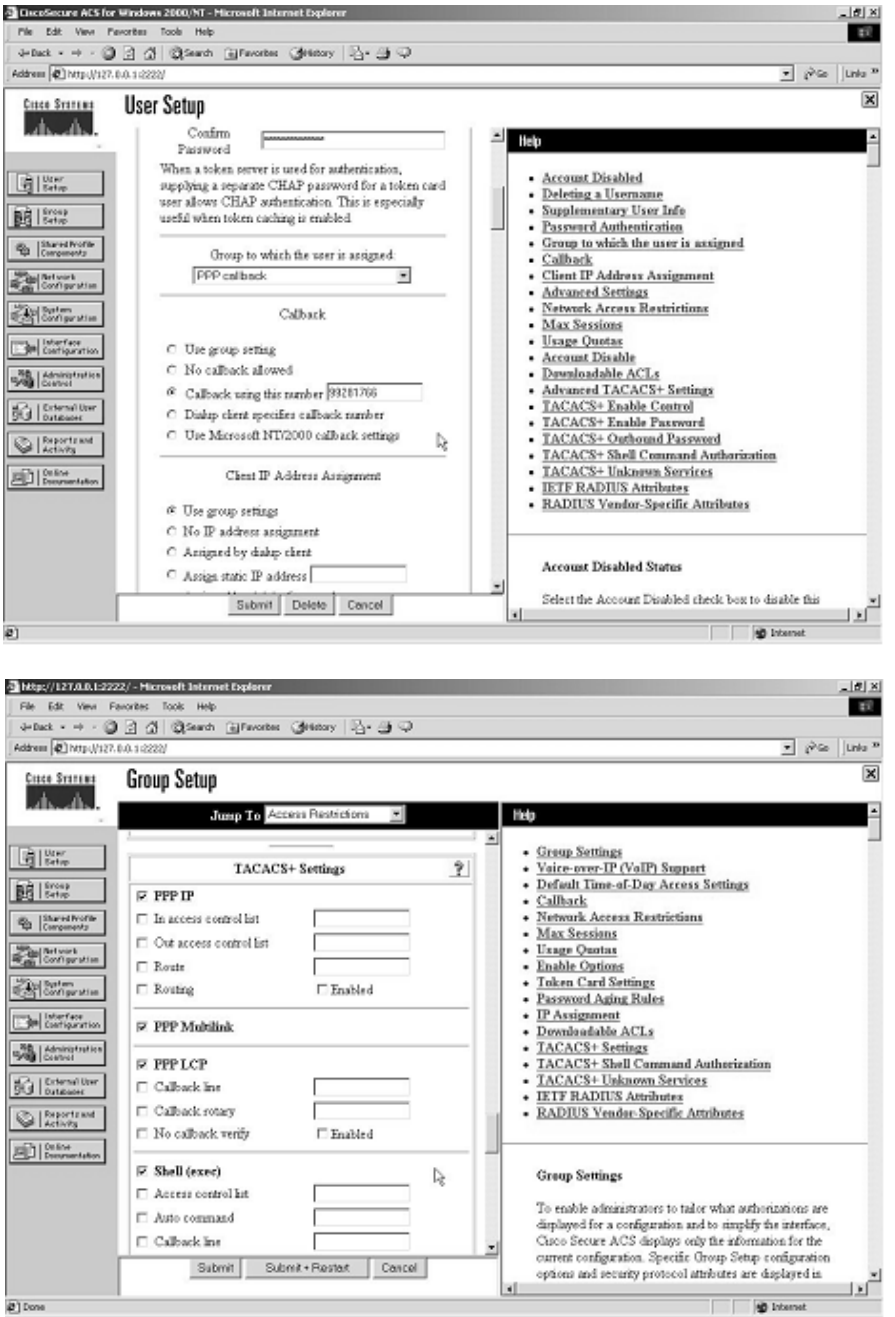

Section 4.0: PIX Configuration

4.1: Basic PIX Configuration

1 As stated earlier, do not configure a default route on PIX. It should learn it from R3 via RIP. Make sure you are able to ping all parts of the network including behind PIX.

4.2: Network Address Translation (NAT)

1 Configure a static NAT on PIX for the syslog server behind PIX.

2 Configure outside access list to open TCP port 1468 for TCP-based reliable syslog server:

```
static (inside,outside) 10.50.31.65 192.168.6.65 netmask 255.255.255.255 0 0
access-list outside permit tcp any host 10.50.31.65 eq 1468 (hitcnt=0)
```

4.3: Advanced Configuration

1 The problem is that PIX is replying for ARP request for the server mentioned. This could be due to a global or alias configured for the same IP address. The fix is to turn off proxy-arp for this interface. **sysopt noproxyarp inside** stops PIX answering for the ARP requests coming from the inside interface.

Section 5.0: IPSec/GRE Configuration

5.1: IPSec

5.1.1: IPSec LAN-to-LAN Using Preshared

1 Configure a LAN-to-LAN IPSec between the PIX and R4. The key is the interesting traffic for IPSec—the IPSec access list, which should be for UDP port 45000, the postoffice protocol communication between the IDS and Director. You can also configure an access list for UDP traffic from host to host—10.50.13.82 to 10.50.31.60.

5.1.2: Advanced IPSec LAN-to-LAN

1 Configure GRE traffic in section 5.2. IPSec access list should be host-to-host and use tunnel mode. Configure ISAKMP keepalive to check the connectivity. If the peer does not respond, phase1 SA will go down and this will also take down the phase 2 SAs.

2 Also remember to configure **no ip route-cache** on all GRE tunnels and physical interfaces where crypto map is applied.

5.2: GRE

1. This is a tricky one. Configure GRE between R3 and R6. You need to configure static translation on PIX for loopback2 to the same address for GRE tunnel on R3 to peer as the GRE destination.

2. Furthermore, modify the outside access list on PIX to allow ESP and UDP/500 from host 133.133.133.133 to 166.166.166.166. You do not need to allow GRE since the packets will be encrypted as per section 5.1.2:

```
access-list outside permit esp host 133.133.133.133 host 166.166.166.166
   (hitcnt=79166)
access-list outside permit udp host 133.133.133.133 host 166.166.166.166
   eq isakmp (hitcnt=99)
static (inside,outside) 166.166.166.166 166.166.166.166
```

3. Redistribute OSPF into EIGRP 100 with a route map to match only loopbacks in area30. The example that follows is for the redistribution configuration on R3:

```
router eigrp 100
  redistribute ospf 110 route-map o2e.
!
access-list 2 permit 30.30.1.0
access-list 2 permit 30.30.2.0
access-list 2 permit 30.30.3.0
access-list 2 permit 30.30.4.0
access-list 2 permit 30.30.5.0
access-list 2 permit 30.30.6.0
!
route-map o2e permit 10
 match ip address 2
```

We are not done yet.

Now, if you do a **show ip route** on R6, you will see that it is learning all the routes via the GRE tunnel interface as expected. See the routing table on R6:

```
r6#show ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route

Gateway of last resort is 10.10.6.1 to network 0.0.0.0

D EX    30.30.1.0 [170/297270016] via 36.36.36.3, 00:13:47, Tunnel63
D EX    30.30.2.0 [170/297270016] via 36.36.36.3, 00:13:47, Tunnel63
D EX    30.30.3.0 [170/297270016] via 36.36.36.3, 00:13:47, Tunnel63
```

```
D EX    30.30.4.0 [170/297270016] via 36.36.36.3, 00:13:47, Tunnel63
D EX    30.30.5.0 [170/297270016] via 36.36.36.3, 00:13:47, Tunnel63
D EX    30.30.6.0 [170/297270016] via 36.36.36.3, 00:13:47, Tunnel63
R*   0.0.0.0/0 [120/1] via 10.10.6.1, 00:00:16, Ethernet0/1
```

The question requires pinging the even networks via the tunnel and odd networks via the PIX. In doing so, it is allowed to use an ACL with one line only.

The solution is "policy routing." You need to create a policy route to match the odd networks and set the next hop to the PIX inside interface—that is, 10.10.6.1—and apply it in global mode, as packets will originate from R6 when testing:

```
ip local policy route-map next-hop
!
access-list 102 permit ip any 30.30.1.0 0.0.254.255
!
route-map next-hop permit 10
  match ip address 102
  set ip next-hop 10.10.6.1
```

The way to confirm if it is working is to turn on **debug icmp trace** on PIX, and ping the odd networks. You will see the packets flowing through, but when you ping even networks, it won't show, as they will be traversing as GRE/IPSec packets and not ICMP traffic. The following example demonstrates this procedure.

```
pix# debug icmp trace
ICMP trace on
Warning: this may cause problems on busy networks

r6#ping 30.30.1.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 30.30.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 104/179/200 ms

r6#ping 30.30.2.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 30.30.2.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 8/8/8 ms

pix# debug icmp trace
ICMP trace on
Warning: this may cause problems on busy networks
pix# 190: Outbound ICMP echo request (len 72 id 58378 seq 1947) 36.36.36.6 >
   10.50.31.1 >30.30.1.1
191: Inbound  ICMP echo reply (len 72 id 17920 seq 1947) 30.30.1.1 >
   10.50.31.1 > 36.36.36.6
192: Outbound ICMP echo request (len 72 id 58634 seq 1947) 36.36.36.6 >
   10.50.31.1 > 30.30.1.1
```

```
193: Inbound  ICMP echo reply (len 72 id 18176 seq 1947) 30.30.1.1 >
   10.50.31.1 > 36.36.36.6
194: Outbound ICMP echo request (len 72 id 58890 seq 1947) 36.36.36.6 >
   10.50.31.1 > 30.30.1.1
195: Inbound  ICMP echo reply (len 72 id 18432 seq 1947) 30.30.1.1 >
   10.50.31.1 > 36.36.36.6
196: Outbound ICMP echo request (len 72 id 59146 seq 1947) 36.36.36.6 >
   10.50.31.1 > 30.30.1.1
197: Inbound  ICMP echo reply (len 72 id 18688 seq 1947) 30.30.1.1 >
   10.50.31.1 > 36.36.36.6
198: Outbound ICMP echo request (len 72 id 59402 seq 1947) 36.36.36.6 >
   10.50.31.1 > 30.30.1.1
199: Inbound  ICMP echo reply (len 72 id 18944 seq 1947) 30.30.1.1 >
   10.50.31.1 > 36.36.36.6
```

As you can see, that ping works for both the odd and even networks from R6, but ICMP packets are only seen when pinging the *odd* network, not the *even* network.

Note that the return echo-reply packets are sent back to the 36.36.36.6 IP address, which is the tunnel IP on R6.

You need to create a static route on PIX for this network, or you will notice that the odd network pings are unsuccessful.

Section 6.0: IOS Firewall Configuration

6.1: CBAC

6.1.1: Basic CBAC Configuration

1. Configure basic IOS Firewall **ip inspect** commands and inspect TCP/UDP/HTTP only. Apply inspect outbound on serial links and ingress ACL for filtering.

6.1.2: Firewall Filtering

1. Inbound ACL on serial links, permit ICMP, OSPF, BGP, and replies from TACACS+ server and host 111.111.111.111 to be able to Telnet to R2.
2. For anti-spoofing, do a **show ip route connected**. Whichever networks are listed should be denied in the ACL for source network:

```
r2#show access-lists 120
Extended IP access list 120
    deny ip 12.12.12.0 0.0.0.255 any
    deny ip 122.122.122.0 0.0.0.255 any
    deny ip 10.50.22.0 0.0.0.15 any
    permit ospf any any (73740 matches)
    permit tcp any any eq bgp (29682 matches)
```

```
permit tcp any eq bgp any (5155 matches)
permit icmp any any (314 matches)
permit tcp host 10.50.31.6 eq tacacs any (100 matches)
permit tcp host 111.111.111.111 any eq telnet (636 matches)
```

6.1.3: Advanced CBAC Configuration

1 Configure TCP embryonic (half-open) connections as follows:

```
ip inspect tcp max-incomplete host 200 block-time 0
```

6.2: Intrusion Detection System (IDS)

6.2.1: Basic IDS Configuration

1 Configure basic IDS on R4 using the **ip audit** command set. Use the first example that follows to configure IDS, and use the second example for logs generated when you detect an attack/signature.

NOTE Note that communication between IDS and Director is on UDP port 45000.

```
ip audit name lab1 info action alarm
ip audit name lab1 attack action alarm
!
interface FastEthernet2/0
 ip address 10.10.45.4 255.255.255.0
 ip audit lab1 in
 ip audit lab1 out
 duplex half
```

```
6d23h: %IDS-4-ICMP_FRAGMENT_SIG: Sig:2150:Fragmented ICMP Traffic - from
   10.10.45.5 to 10.10.45.4
```

6.2.2: Signature Tuning

1 If you receive false positive alarms from the IDS on R4, you need to disable signature 3050 for host 10.50.16.5 on R4. The following example demonstrates tuning IDS signatures on R4:

```
ip audit signature 3050 list 5
!
access-list 5 deny   10.50.16.5
access-list 5 permit any
```

6.2.3: Spam Attack

1 Configure R4 protection against SMTP mail spamming using the following command:

```
ip audit smtp spam 500
```

Section 7.0: AAA

7.1: AAA on the Router

1 Configure AAA on R4 to use the TACACS+ server.

2 Configure authentication, EXEC authorization, and command-level 1/10/15 authorization.

3 Move the **show running-config** command to level 10 for user1 to be able to invoke it.

4 Configure fallback to local in the event the AAA server goes down.

5 Make sure you use a named method list and apply it to vty lines. Do not configure any authentication or authorization for console or auxiliary ports, or you will lose all marks.

6 Use the following example to configure all of the above.

```
aaa new-model
aaa authentication login vtyline group tacacs+ local
aaa authentication login con-none none
aaa authorization exec vtyexec group tacacs+ local
aaa authorization exec conexec none
aaa authorization commands 1 comm1 group tacacs+ local
aaa authorization commands 1 comm-con-none none
aaa authorization commands 10 comm10 group tacacs+ local
aaa authorization commands 10 comm-con-none none
aaa authorization commands 15 comm15 group tacacs+ local
aaa authorization commands 15 comm-con-none none
!
username user1 privilege 10 password 7 044E18031D70
username user2 privilege 15 password 7 13100417195E
!
privilege exec level 10 show run
privilege exec level 15 show!
line con 0
 exec-timeout 0 0
 authorization commands 1 comm-con-none
 authorization commands 10 comm-con-none
 authorization commands 15 comm-con-none
 authorization exec conexec
 login authentication con-none
line aux 0
 authorization commands 1 comm-con-none
```

```
 authorization commands 10 comm-con-none
 authorization commands 15 comm-con-none
 authorization exec conexec
 login authentication con-none
line vty 0 4
 authorization commands 1 comm1
 authorization commands 10 comm10
 authorization commands 15 comm15
 authorization exec vtyexec
 login authentication vtyline
!
end
```

7 Configure ACS with two users as follows.

User1 with privilege level 10 and permit the **show run** command. See Figure 1-6 for user settings on CiscoSecure ACS.

User2 with privilege level 15 with all commands permitted. See Figure 1-7 for user settings on CiscoSecure ACS.

8 Configure CiscoSecure ACS users above with corresponding privilege levels, so when they log in, they land in enable mode and don't need to enter **enable**. You need to configure exec authorization to achieve this task. Refer to Figure 1-6 for user1 and Figure 1-7 for user2 profile settings on ACS.

Figure 1-6 *User1 Settings on CiscoSecure ACS*

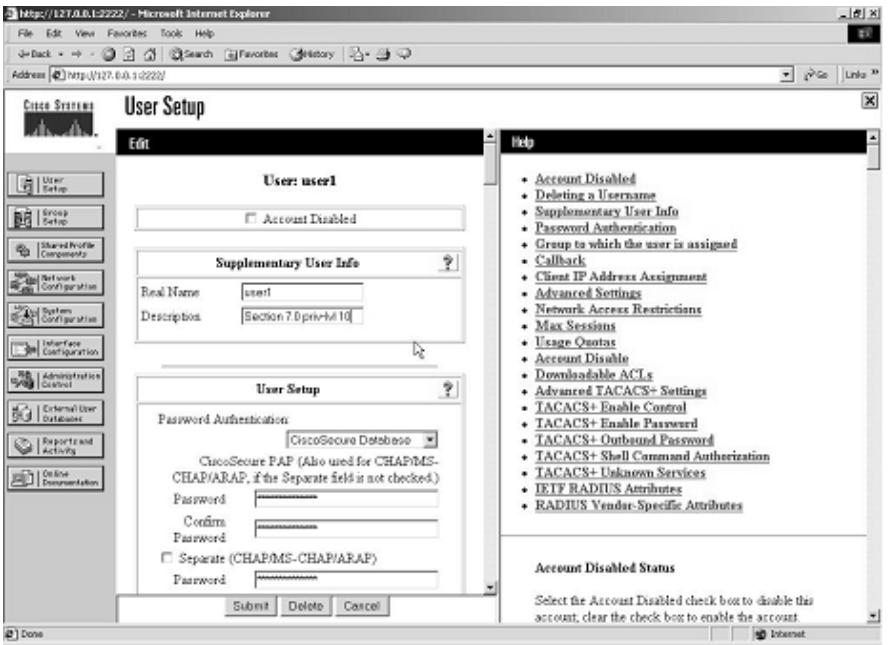

Figure 1-6 *User1 Settings on CiscoSecure ACS (Continued)*

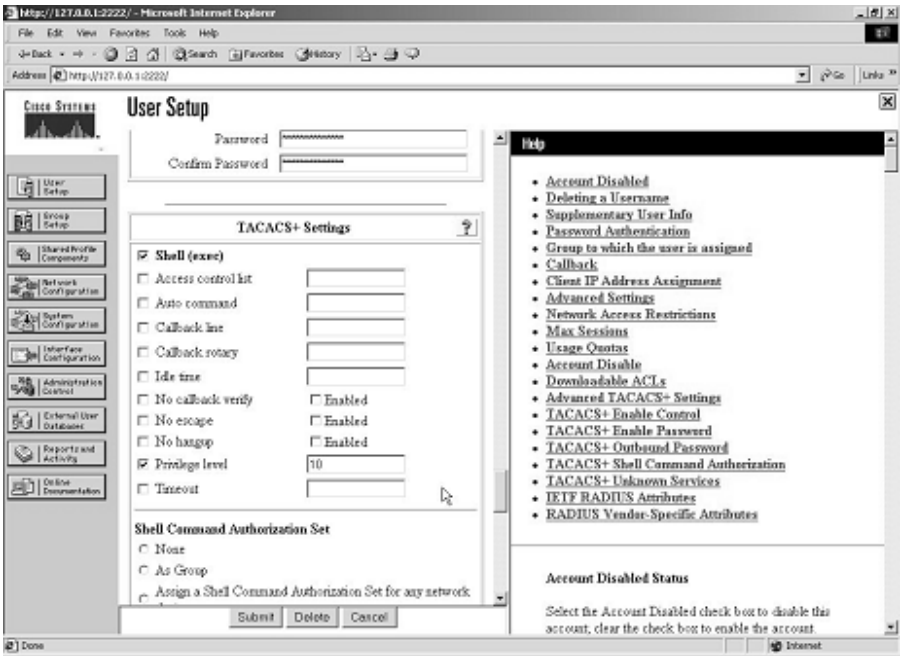

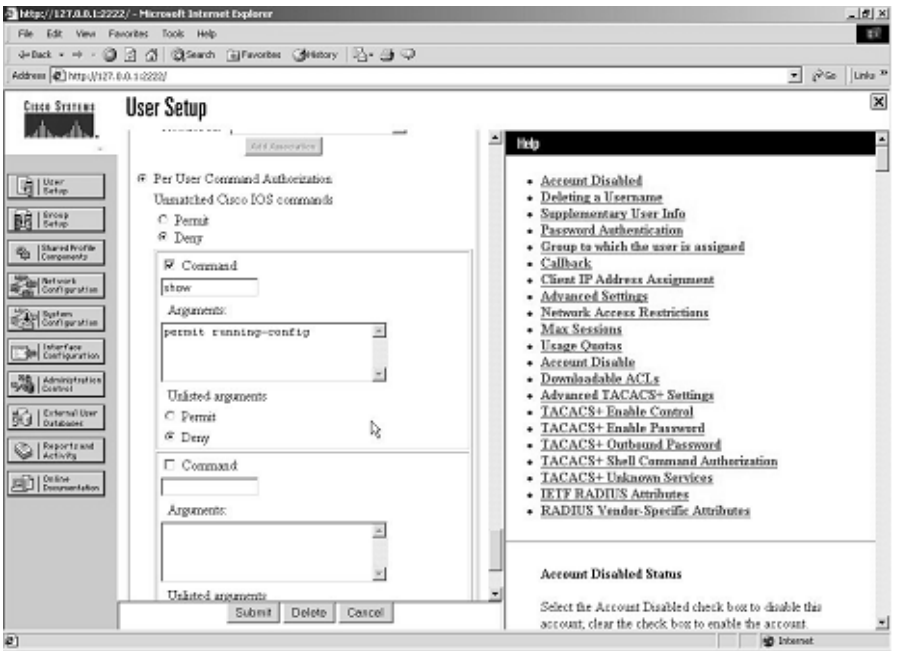

Figure 1-7 *User2 Settings on CiscoSecure ACS*

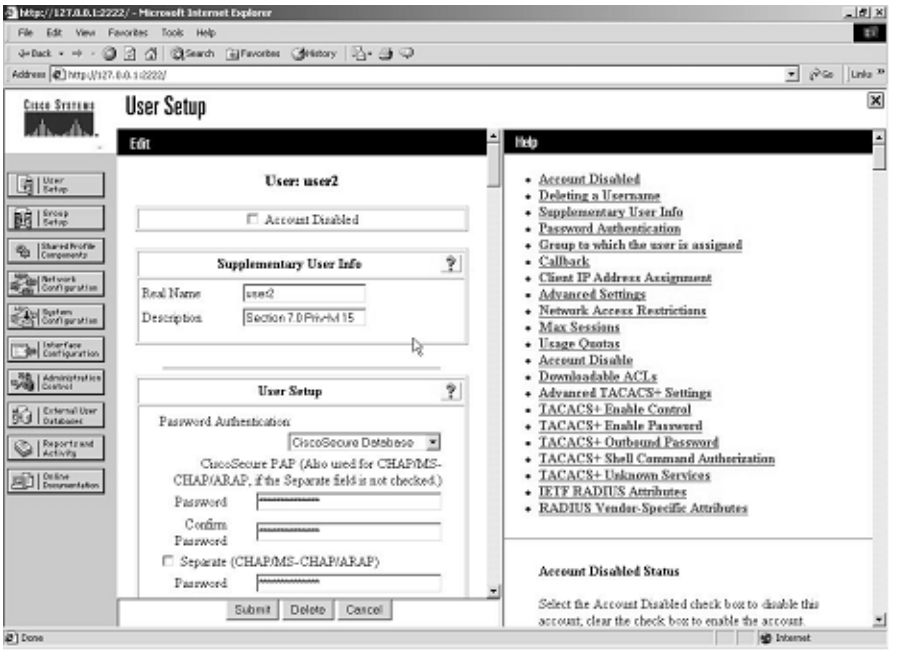

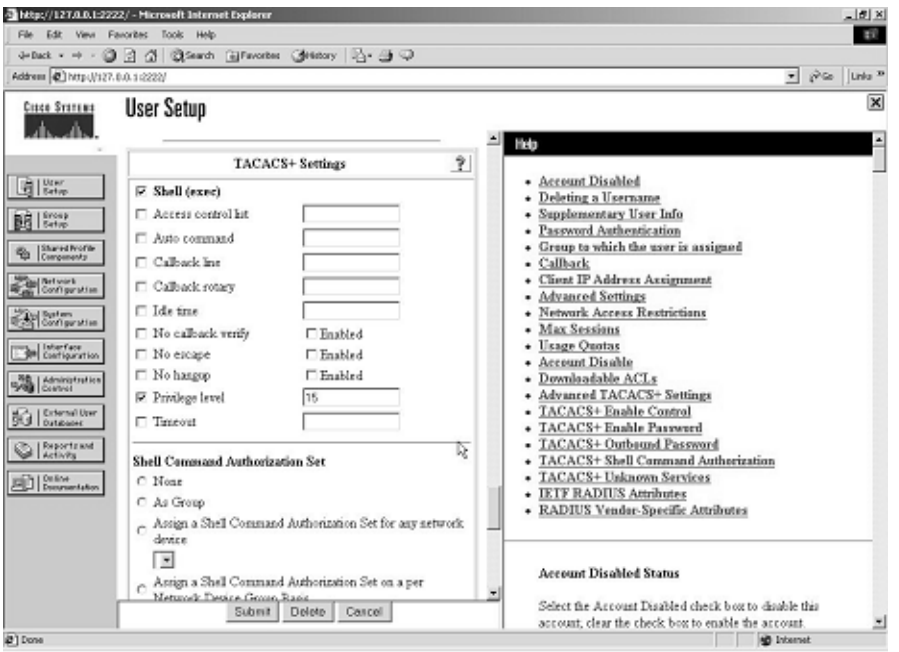

Figure 1-7 *User2 Settings on CiscoSecure ACS (Continued)*

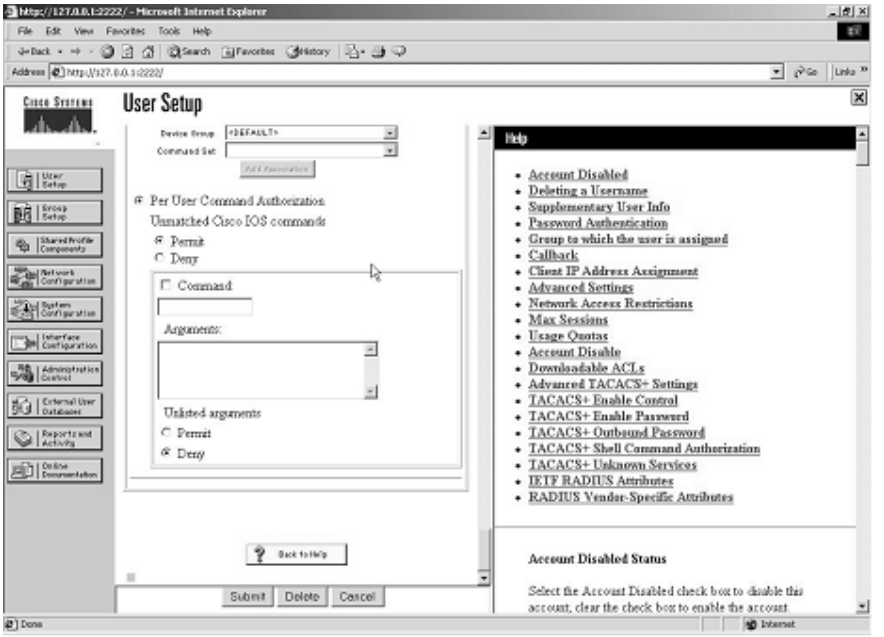

7.2: AAA on PIX

1. Configure TACACS+ authentication and authorization for Telnet service on PIX (refer to the example that follows item 3).
2. Configure static translation for Loopback1 of R6. (Refer to the example that follows item 3 to configure the PIX.)
3. Configure username **r6telnet** on ACS with Per User Command Authorization set to permit Telnet service for R6 Loopback1 only. Refer to Figure 1-8 for **r6telnet** profile settings on ACS.

```
pix# show aaa
aaa authentication include telnet outside 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0 ACS
aaa authorization include telnet outside 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0 ACS
pix#
pix# show aaa-server
aaa-server ACS (inside) host 192.168.6.6 cisco timeout 10
pix#
pix(config)# show access-list outside
access-list outside permit tcp any host 10.50.31.6 eq tacacs (hitcnt=103)
access-list outside permit tcp any host 16.16.16.16 eq telnet (hitcnt=7)
pix(config)# show static
static (inside,outside) 16.16.16.16 16.16.16.16 netmask 255.255.255.255 0 0
```

```
! Login capture from R3 telnetting to R6 loopback1:
r3#telnet 16.16.16.16
Trying 16.16.16.16 ... Open

Username: r6telnet
Password: r6telnet

User Access Verification
Password:
r6>en
Password:
r6#
r6#
! After successfully logging on to R6, confirm that
! authentication/authorization is working on pix;
pix# show uauth
                        Current    Most Seen
Authenticated Users       1          1
Authen In Progress        0          1
user 'r6telnet' at 10.50.31.2, authorized to:
   port 16.16.16.16/telnet
   absolute   timeout: 0:05:00
   inactivity timeout: 0:00:00
```

Figure 1-8 **r6telnet** *Settings on CiscoSecure ACS*

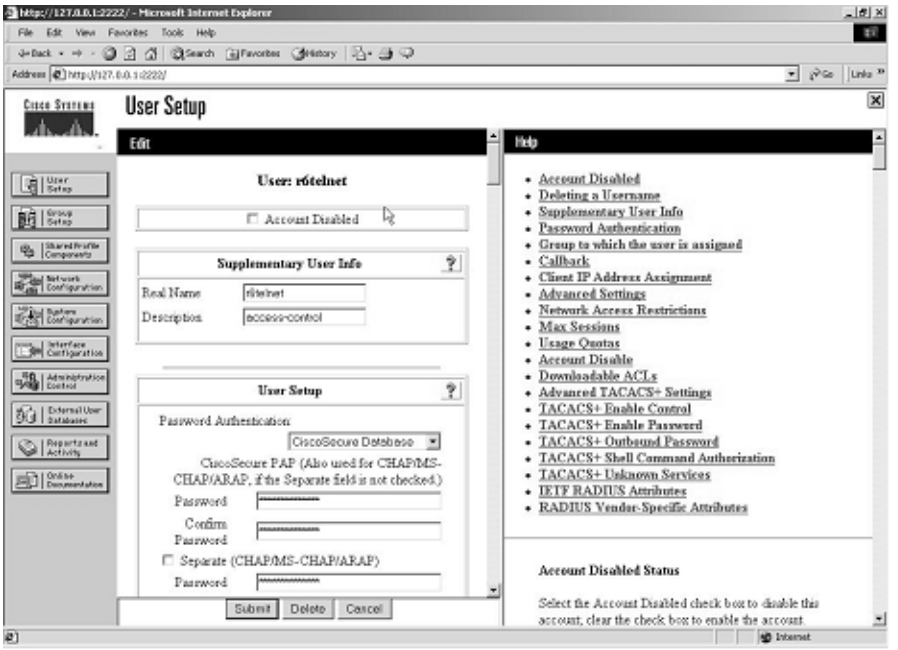

Figure 1-8 **r6telnet** *Settings on CiscoSecure ACS (Continued)*

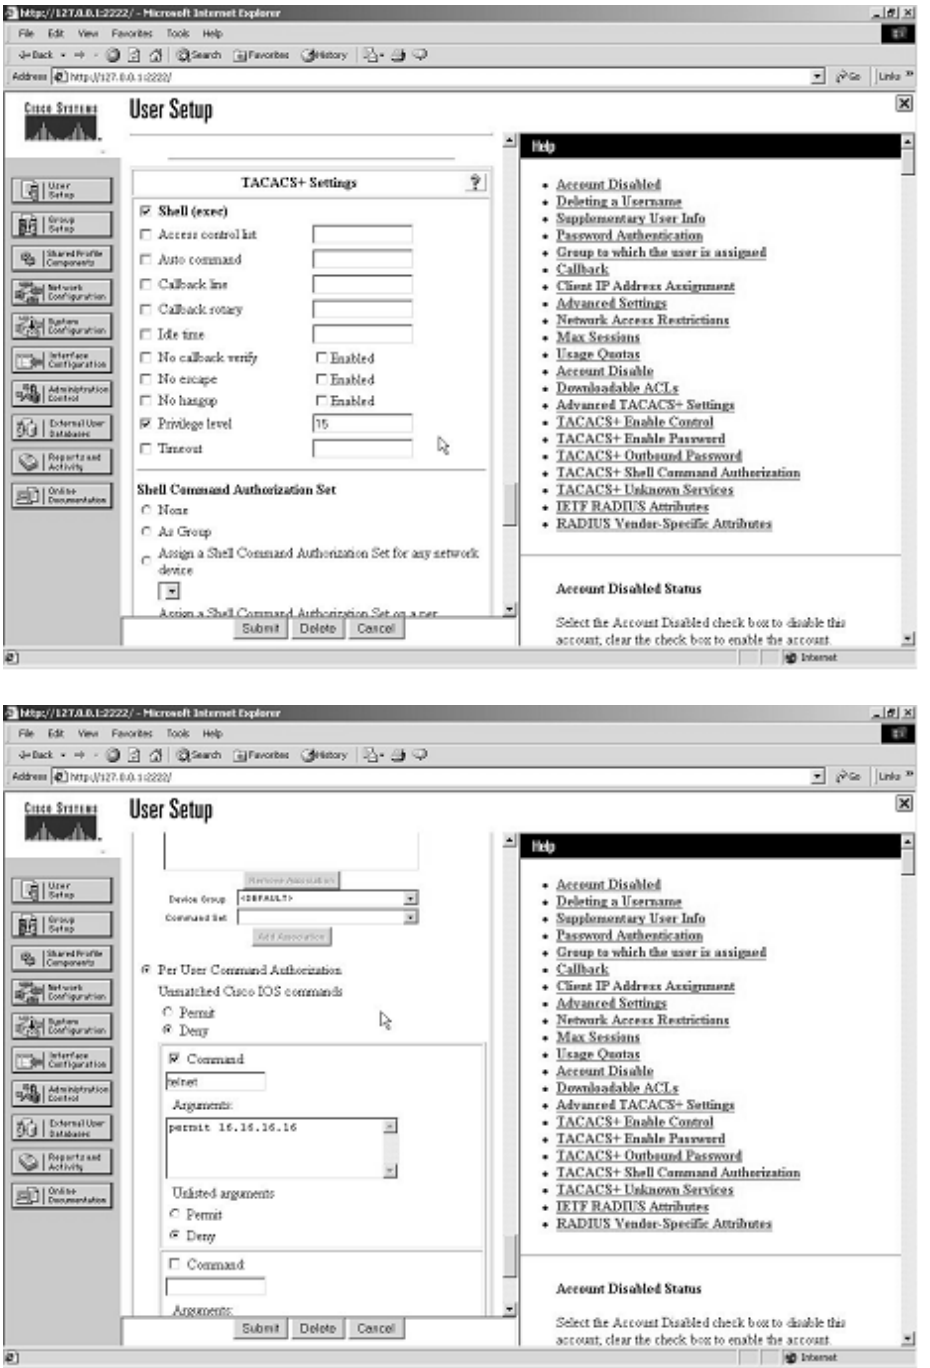

NOTE If Shell Command Authorization Set does not appear in User Setup in ACS, go to Interface Configuration and select TACACS+ and tick the User column for Shell (exec). See Figure 1-9.

Figure 1-9 *Interface Configuration on ACS*

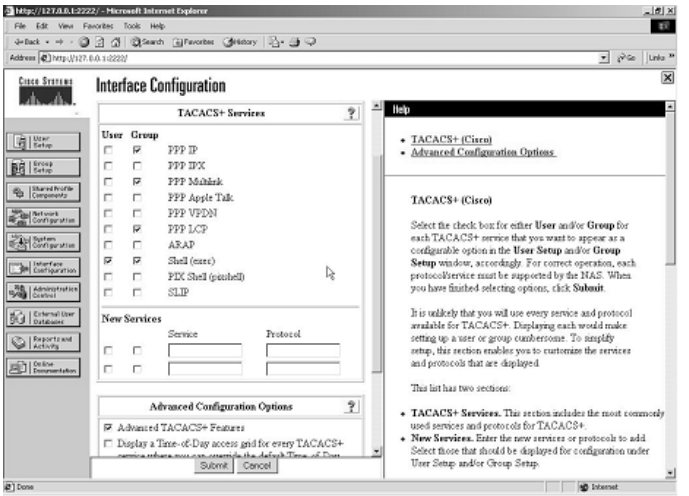

NOTE The Reports and Activity section in CiscoSecure ACS is very useful for troubleshooting. Verify FAILED/PASSED attempts in Reports, as shown in Figure 1-10.

Figure 1-10 *Reports and Activity in ACS*

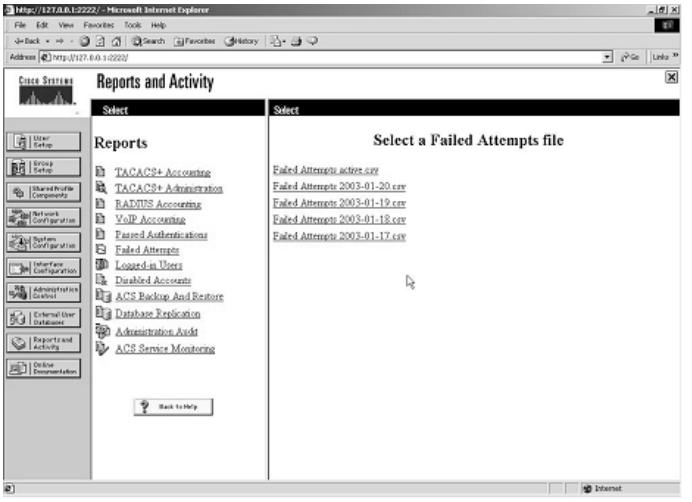

Section 8.0: Advanced Security

8.1: Password Protection

1 Configure **service password-encryption** on all the routers to encrypt the enable password; otherwise, they will appear in clear text in the configuration.

8.2: EXEC Authentication

1 Configure **enable secret** on R2.

2 Configure authentication for shell EXEC without using the AAA engine using the **enable use-tacacs** command. Note that this is not TACACS+ but TACACS server (without the +). CiscoSecure ACS is not a TACACS server but TACACS+ only.

3 Configure fallback to pass authentication in the event the TACACS server is down or not found using **enable last-resort succeed**.

8.3: Access Control

1 In this case, you can configure **autocommand** for a user to Telnet to the router. **autocommand** will execute the required command and exit the session. This way the user will not be able to keep its Telnet session:

```
username testconfig privilege 15 password 7 15060E1F1029242A2E3A32
username testconfig autocommand show run
!
line vty 0 4
 privilege level 15
 password 7 110A1016141D
 login local
!
end
```

Test by Telnetting from R1 to 10.50.13.2.

```
r1#telnet 10.50.13.2
Trying 10.50.13.2 ... Open

User Access Verification

Username: testconfig
Password: testconfig
Building configuration...
```

```
Current configuration : 7022 bytes
!
! Last configuration change at 23:46:49 AEDT Sun Jan 19 2003
! NVRAM config last updated at 00:15:25 AEDT Mon Jan 20 2003
!
version 12.1
no service single-slot-reload-enable
service timestamps debug uptime
service timestamps log uptime
service password-encryption
!
hostname r3
!
snipped
!
end

[Connection to 10.50.13.2 closed by foreign host]
r1#
```

As you can see, as soon as the **show run** command output finished, the session was closed.

2 Configure R5 Telnet access to permit host 133.133.133.133 only:

```
access-list 3 permit 133.133.133.133
!
line vty 0 4
 access-class 3 in
 password 7 13061E010803
 login
!
end
```

Section 9.0: IP Services and Protocol-Independent Features

9.1: NAT

1 Configure NAT for Loopback3 192.168.3.1/24.

2 The objective is that when sourced from Loopback3 to anywhere on the network, it should be translated using the egress interface. For example, if you ping 122.122.122.122, it

will use egress interface Serial1/0.3, whereas if you ping 144.144.144.144, it will use egress interface Serial1/0.1. If you ping 166.166.166.166, it will use egress interface FastEthernet0/0. To configure this multihomed NAT, enter the following:

```
ip nat inside source route-map fastethernet0/0 interface FastEthernet0/0
   overload
ip nat inside source route-map serial1/0.1 interface Serial1/0.1 overload
ip nat inside source route-map serial1/0.3 interface Serial1/0.3 overload
!
access-list 102 permit ip 192.168.3.0 0.0.0.255 any
!
route-map serial1/0.1 permit 10
 match ip address 102
 match interface Serial1/0.1
!
route-map serial1/0.3 permit 10
 match ip address 102
 match interface Serial1/0.3
!
route-map fastethernet0/0 permit 10
 match ip address 102
 match interface FastEthernet0/0
```

To test multihomed NAT, enter the following:

```
! "Debug ip nat" on R3 and ping 122.122.122.122, 144.144.144.144 and
   166.166.166.166
! sourcing from Loopback3:
r3#ping ip
Target IP address: 122.122.122.122
Repeat count [5]:
Datagram size [100]:
Timeout in seconds [2]:
Extended commands [n]: y
Source address or interface: loopback3
Type of service [0]:
Set DF bit in IP header? [no]:
Validate reply data? [no]:
Data pattern [0xABCD]:
Loose, Strict, Record, Timestamp, Verbose[none]:
Sweep range of sizes [n]:
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 122.122.122.122, timeout is 2 seconds:
!!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 68/68/68 ms
r3#
r3#
4d14h: NAT: s=192.168.3.1->10.50.13.18, d=122.122.122.122 [195]
4d14h: NAT*: s=122.122.122.122, d=10.50.13.18->192.168.3.1 [195]
4d14h: NAT: s=192.168.3.1->10.50.13.18, d=122.122.122.122 [196]
4d14h: NAT*: s=122.122.122.122, d=10.50.13.18->192.168.3.1 [196]
4d14h: NAT: s=192.168.3.1->10.50.13.18, d=122.122.122.122 [197]
4d14h: NAT*: s=122.122.122.122, d=10.50.13.18->192.168.3.1 [197]
4d14h: NAT: s=192.168.3.1->10.50.13.18, d=122.122.122.122 [198]
4d14h: NAT*: s=122.122.122.122, d=10.50.13.18->192.168.3.1 [198]
4d14h: NAT: s=192.168.3.1->10.50.13.18, d=122.122.122.122 [199]
4d14h: NAT*: s=122.122.122.122, d=10.50.13.18->192.168.3.1 [199]
r3#
r3#
r3#ping ip
Target IP address: 144.144.144.144
Repeat count [5]:
Datagram size [100]:
Timeout in seconds [2]:
Extended commands [n]: y
Source address or interface: loopback3
Type of service [0]:
Set DF bit in IP header? [no]:
Validate reply data? [no]:
Data pattern [0xABCD]:
Loose, Strict, Record, Timestamp, Verbose[none]:
Sweep range of sizes [n]:
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 144.144.144.144, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 96/99/101 ms
r3#
r3#
4d14h: NAT: s=192.168.3.1->10.50.13.2, d=144.144.144.144 [210]
4d14h: NAT*: s=144.144.144.144, d=10.50.13.2->192.168.3.1 [210]
4d14h: NAT: s=192.168.3.1->10.50.13.2, d=144.144.144.144 [211]
4d14h: NAT*: s=144.144.144.144, d=10.50.13.2->192.168.3.1 [211]
4d14h: NAT: s=192.168.3.1->10.50.13.2, d=144.144.144.144 [212]
4d14h: NAT*: s=144.144.144.144, d=10.50.13.2->192.168.3.1 [212]
4d14h: NAT: s=192.168.3.1->10.50.13.2, d=144.144.144.144 [213]
4d14h: NAT*: s=144.144.144.144, d=10.50.13.2->192.168.3.1 [213]
```

```
4d14h: NAT: s=192.168.3.1->10.50.13.2, d=144.144.144.144 [214]
4d14h: NAT*: s=144.144.144.144, d=10.50.13.2->192.168.3.1 [214]
r3#
r3#
r3#ping ip
Target IP address: 166.166.166.166
Repeat count [5]:
Datagram size [100]:
Timeout in seconds [2]:
Extended commands [n]: y
Source address or interface: loopback3
Type of service [0]:
Set DF bit in IP header? [no]:
Validate reply data? [no]:
Data pattern [0xABCD]:
Loose, Strict, Record, Timestamp, Verbose[none]:
Sweep range of sizes [n]:
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 166.166.166.166, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 4/4/4 ms
r3#
r3#
4d14h: NAT: s=192.168.3.1->10.50.31.2, d=166.166.166.166 [205]
4d14h: NAT*: s=166.166.166.166, d=10.50.31.2->192.168.3.1 [205]
4d14h: NAT: s=192.168.3.1->10.50.31.2, d=166.166.166.166 [206]
4d14h: NAT*: s=166.166.166.166, d=10.50.31.2->192.168.3.1 [206]
4d14h: NAT: s=192.168.3.1->10.50.31.2, d=166.166.166.166 [207]
4d14h: NAT*: s=166.166.166.166, d=10.50.31.2->192.168.3.1 [207]
4d14h: NAT: s=192.168.3.1->10.50.31.2, d=166.166.166.166 [208]
4d14h: NAT*: s=166.166.166.166, d=10.50.31.2->192.168.3.1 [208]
4d14h: NAT: s=192.168.3.1->10.50.31.2, d=166.166.166.166 [209]
4d14h: NAT*: s=166.166.166.166, d=10.50.31.2->192.168.3.1 [209]
```

The preceding test from R3 confirms NATing loopback3 with respective egress interface as per the route map:

Ping 122.122.122.122 NATed with 10.50.13.18 egress Serial1/0.3

Ping 144.144.144.144 NATed with 10.50.13.2 egress Serial1/0.1

Ping 166.166.166.166 NATed with 10.50.31.2 egress FastEthernet0/0

9.2: NTP

1 Configure R2 as NTP Server and R1 as NTP Client.

2 Configure authentication using the md5 key. NTP status and authentication on R2 is as follows:

```
r1# show ntp status
Clock is synchronized, stratum 9, reference is 10.50.13.34
nominal freq is 250.0000 Hz, actual freq is 250.0000 Hz, precision is 2**24
reference time is C1D5BFAA.20689871 (00:22:02.126 UTC Mon Jan 20 2003)
clock offset is 1.6778 msec, root delay is 64.39 msec
root dispersion is 126.82 msec, peer dispersion is 0.12 msec
r1#
r1#
r1#show ntp associations detail
10.50.13.34 configured, authenticated, our_master, sane, valid, stratum 8
ref ID 127.127.7.1, time C1D5BF88.FE740124 (00:21:28.993 UTC Mon Jan 20 2003)
our mode client, peer mode server, our poll intvl 64, peer poll intvl 64
root delay 0.00 msec, root disp 125.03, reach 377, sync dist 157.349
delay 64.39 msec, offset 1.6778 msec, dispersion 0.12
precision 2**16, version 3
org time C1D5BFAA.188E6A78 (00:22:02.095 UTC Mon Jan 20 2003)
rcv time C1D5BFAA.20689871 (00:22:02.126 UTC Mon Jan 20 2003)
xmt time C1D5BFAA.0FC3685F (00:22:02.061 UTC Mon Jan 20 2003)
filtdelay =    64.67   64.39   64.50   64.45   64.67   64.39   64.80   67.99
filtoffset =    1.66    1.68    1.60    1.55    1.57    1.55    1.66   -0.13
filterror =     0.02    0.03    0.05    0.06    0.08    0.09    0.11    0.12
r1#
r1#
r1#show clock
00:25:19.586 UTC Mon Jan 20 2003
r1#
r1#
```

3 In some IOS it is necessary to enter the NTP authentication commands in a particular order. Below is the exact order that confirms operation:

For R2 (master) enter commands in the following sequence:

```
ntp authentication-key 1 md5 cisco
ntp master 2
```

For R1 (Client) enter commands in the following sequence:

```
ntp authentication-key 1 md5 cisco
ntp authenticate
```

```
ntp trusted-key 1
ntp server 10.50.13.34 key 1
```

4 Remember that you have an inbound access list applied to the serial link on R2; you need to allow NTP.

9.3: SNMP

1 Configure R3 to send SNMP traps when a configuration change happens for BGP:

```
snmp-server community public RO
snmp-server community private RW
snmp-server enable traps config
snmp-server enable traps bgp
snmp-server host 10.50.31.99 public  config bgp

! snip from R3 test using debug snmp packet;

r3#debug snmp packets
SNMP packet debugging is on
r3#
r3#config terminal
Enter configuration commands, one per line.  End with CNTL/Z.
r3(config)#
r3(config)#
5d00h: SNMP: Queuing packet to 10.50.31.99
5d00h: SNMP: V1 Trap, ent ciscoConfigManMIB.2, addr 10.50.31.2, gentrap 6,
   spectrap 1
 ccmHistoryEventEntry.3.162 = 1
 ccmHistoryEventEntry.4.162 = 2
 ccmHistoryEventEntry.5.162 = 3
5d00h: SNMP: Packet sent via UDP to 10.50.31.99
r3(config)#
r3(config)#end
r3#
r3#clear ip bgp *
r3#
5d00h: %BGP-5-ADJCHANGE: neighbor 10.50.13.1 Down User reset
5d00h: SNMP: Queuing packet to 10.50.31.99
5d00h: SNMP: V1 Trap, ent bgp, addr 10.50.31.2, gentrap 6, spectrap 2
 bgpPeerEntry.14.10.50.13.1 = 00 00
 bgpPeerEntry.2.10.50.13.1 = 1
5d00h: %BGP-5-ADJCHANGE: neighbor 10.50.13.17 Down User reset
```

```
5d00h: SNMP: Queuing packet to 10.50.31.99
5d00h: SNMP: V1 Trap, ent bgp, addr 10.50.31.2, gentrap 6, spectrap 2
 bgpPeerEntry.14.10.50.13.17 = 00 00
 bgpPeerEntry.2.10.50.13.17 = 1
5d00h: %BGP-5-ADJCHANGE: neighbor 10.50.31.22 Down User reset
r3#
5d00h: SNMP: Queuing packet to 10.50.31.99
5d00h: SNMP: V1 Trap, ent bgp, addr 10.50.31.2, gentrap 6, spectrap 2
 bgpPeerEntry.14.10.50.31.22 = 04 00
 bgpPeerEntry.2.10.50.31.22 = 1
5d00h: SNMP: Packet sent via UDP to 10.50.31.99
5d00h: SNMP: Packet sent via UDP to 10.50.31.99
5d00h: SNMP: Packet sent via UDP to 10.50.31.99
r3#
r3#
! Snip from PIX config and ACL;
pix# show access-list outside
access-list outside permit udp host 10.50.31.2 host 10.50.31.99 eq snmptrap
   (hitcnt=44)
pix# show static
static (inside,outside) 10.50.31.99 192.168.6.99 netmask 255.255.255.255 0 0
pix#
```

9.4: Policy Routing

1 Configure policy routing on R1 to change the next hop for mail and web server off R3:

```
interface Serial2/0.2 point-to-point
 ip address 10.50.13.33 255.255.255.240
 ip policy route-map server
!
interface Serial2/0.3 point-to-point
 ip address 10.50.13.1 255.255.255.240
 ip policy route-map server
!
!
ip local policy route-map server
!
access-list 101 permit ip any host 10.50.31.98
access-list 102 permit ip any host 10.50.31.99
!
route-map server permit 10
```

```
 match ip address 101
 set ip next-hop 10.50.13.34
!
route-map server permit 20
 match ip address 102
 set ip next-hop 10.50.13.2
!
route-map server permit 30


! Verify with traceroute;
r1#traceroute 10.50.31.98
Type escape sequence to abort.
Tracing the route to 10.50.31.98

  1 10.50.13.34 !A  *  !A


r1#traceroute 10.50.31.99
Type escape sequence to abort.
Tracing the route to 10.50.31.99

  1 10.50.13.2 32 msec 32 msec 32 msec
  2  *  *  *
```

Section 10.0: Security Violations

10.1: Denial of Service—DoS

1 Configure CAR (rate-limit) on R3 to prevent ICMP flooding:

```
interface Serial1/0.1 point-to-point
 ip address 10.50.13.2 255.255.255.240
rate-limit input access-group 110 560000 256000 384000 conform-action
   continue exceed-action drop
!
interface Serial1/0.3 point-to-point
 ip address 10.50.13.18 255.255.255.240
rate-limit input access-group 110 560000 256000 384000 conform-action
   continue exceed-action drop
    !
access-list 110 permit icmp any any
```

10.2: IP Spoofing

1 Configure Unicast RPF IP spoofing protection on PIX for inside and outside interfaces:

```
pix# show ip verify
ip verify reverse-path interface outside
ip verify reverse-path interface inside
```