CCIE

Preparing for the CCIE Routing and Switching Lab

Version 2.1

Student Guide

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the Cisco Web site at www.cisco.com/go/offices.

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia • Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary India • Indonesia • Ireland • Israel • Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden • Switzerland • Taiwan • Thailand • Turkey Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

CCIE e-Prep – Preparing for the CCIE Routing and Switching Lab

MODULE 6: LINK-STATE ROUTING PROTOCOLS

MODULE 7: BGP TECHNOLOGIES

1

Course Introduction

Overview

The Cisco Certified Internetworking Expert (CCIE) Prep course helps qualified CCIE candidates prepare for the Hands-on Lab Exam.

Major topics covered include Frame Relay, Integrated Services Digital Network (ISDN), Asynchronous Transfer Mode (ATM), Layer 2 Switching, Routing Protocols, Desktop Protocols, Multicasting, Voice over Internet Protocol (VoIP), Quality of Service (QoS), and Security.

Outline

The Course Introduction includes these topics:

- Course Objectives
- Cisco's Certification Track
- Learner Skills and Knowledge
- Learner Responsibilities
- General Administration
- Course Roadmap
- Icons and Symbols
- Learner Introductions
- Lab Registration
- What to Expect the Day of the Lab

Course Objectives

This section lists the course objectives.

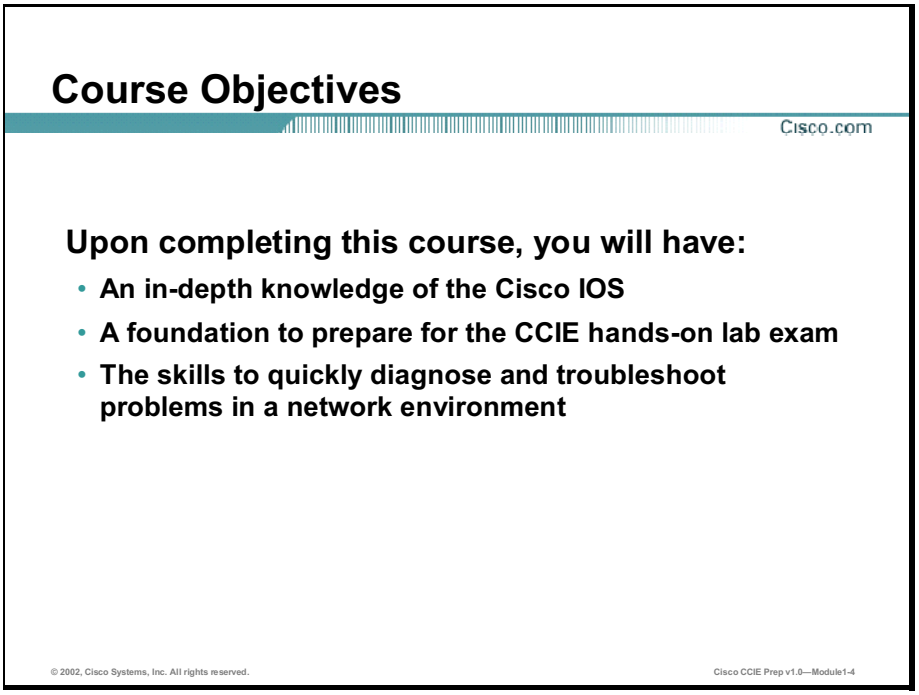

Upon completing this course, you will have:

- An in-depth knowledge of the Cisco Internetwork Operating System (IOS)
- A foundation to prepare for the CCIE hands-on lab exam
- The skills to quickly diagnose and troubleshoot problems in a network environment

Cisco's Certification Track

This section lists the certification requirements of this course.

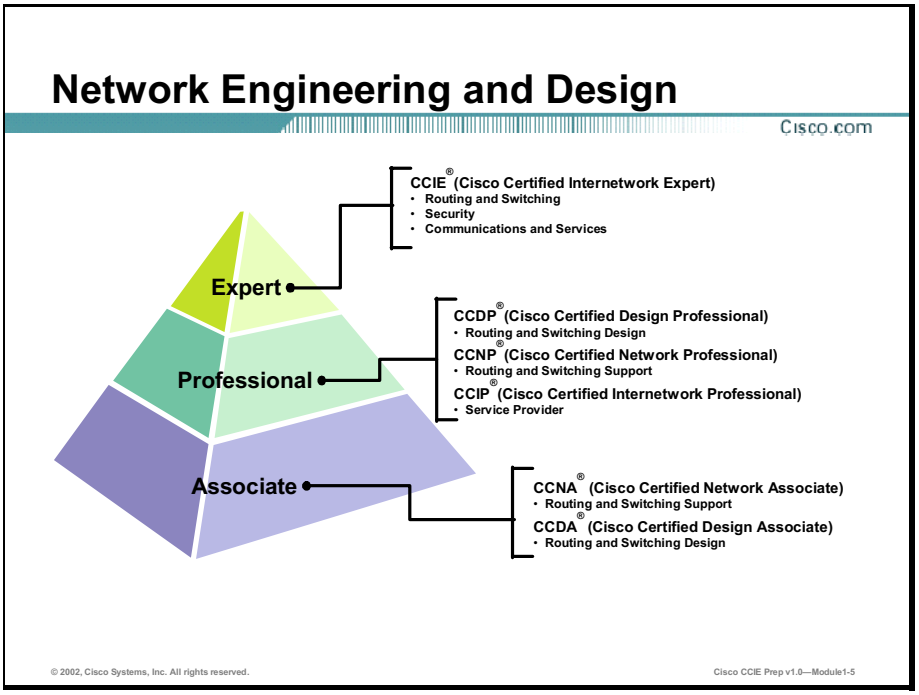

The CCIE program helps individuals, companies, industries, and countries succeed in the networked world by distinguishing the top echelon of internetworking experts.

The program identifies leaders with a proven commitment to their career, the industry, and the process of ongoing learning. While individuals inevitably gain extensive product knowledge on their way to certification, product training is not the CCIE program objective. Rather, the focus is on identifying those experts capable of understanding and navigating the subtleties, intricacies and potential pitfalls inherent to end-to-end networking regardless of technology or product brand.

Learner Skills and Knowledge

This section lists the course prerequisites.

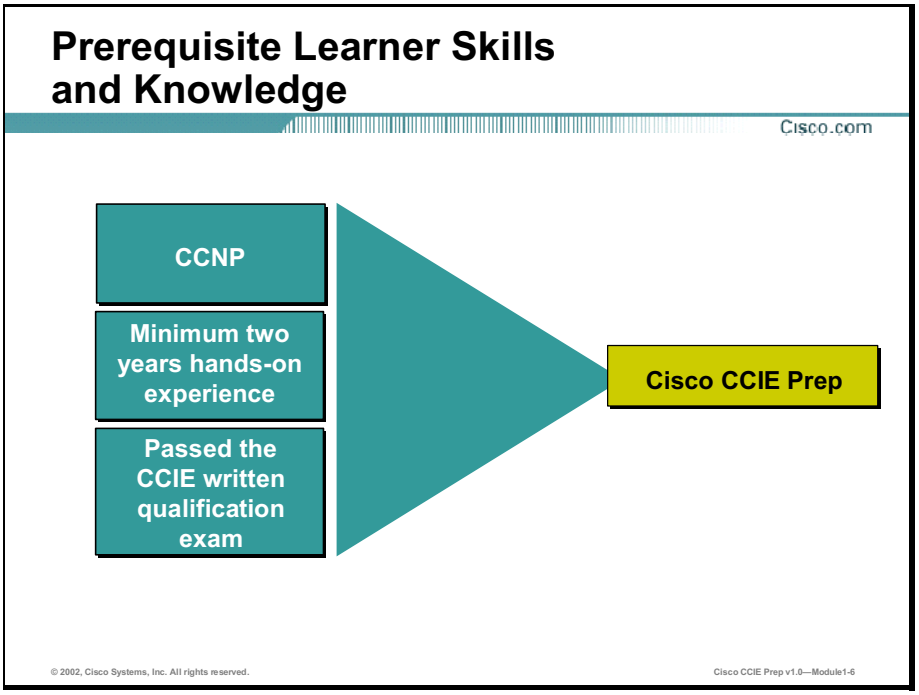

To fully benefit from this course, you must have these prerequisite skills and knowledge:

- Cisco Certified Network Professional (CCNP)
- Minimum two years hands-on experience
- Passed the CCIE written qualification exam

Learner Responsibilities

This section discusses the responsibilities of the learners.

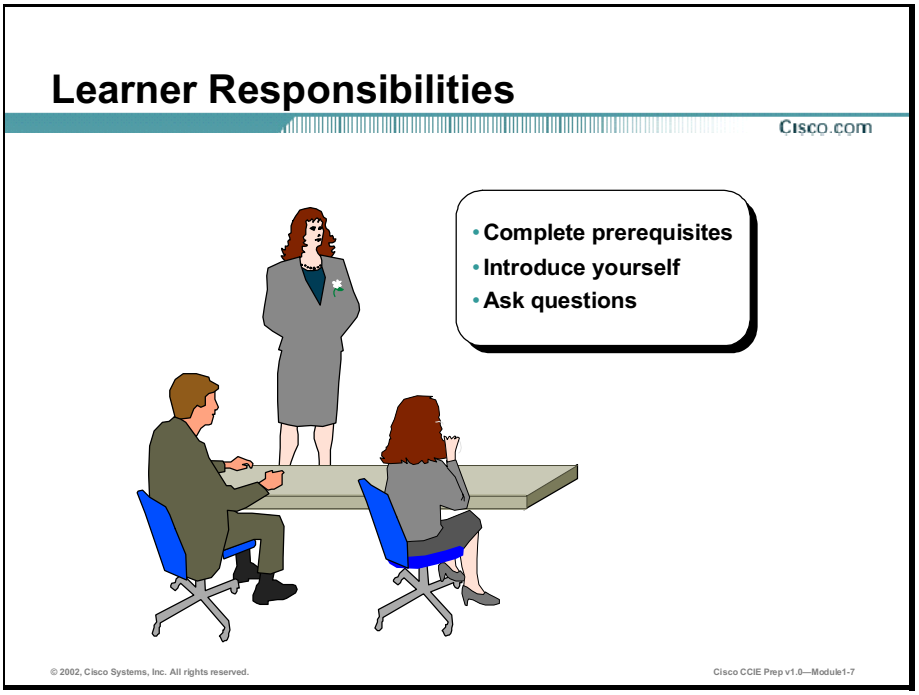

To take full advantage of the information presented in this course, you must have completed the prerequisite requirements.

In class, you are expected to participate in all lesson exercises and assessments.

In addition, you are encouraged to ask any questions relevant to the course materials.

If you have pertinent information or questions concerning future Cisco product releases and product features, please discuss these topics during breaks or after class. The instructor will answer your questions or direct you to an appropriate information source.

General Administration

This section lists the administrative issues for the course.

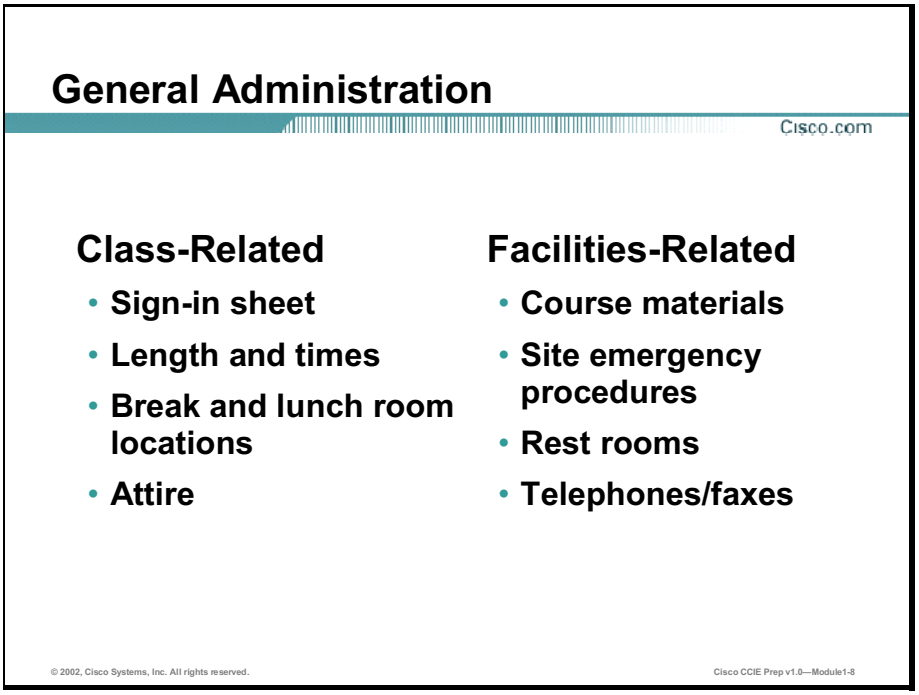

The instructor will discuss the administrative issues noted here so you know exactly what to expect from the class.

- Sign-in process
- Starting and anticipated ending times of each class day
- Class breaks and lunch facilities
- Appropriate attire during class
- Materials you can expect to receive during class
- What to do in the event of an emergency
- Location of the rest rooms
- How to send and receive telephone and fax messages

Course Roadmap

This section covers the suggested flow of the course materials.

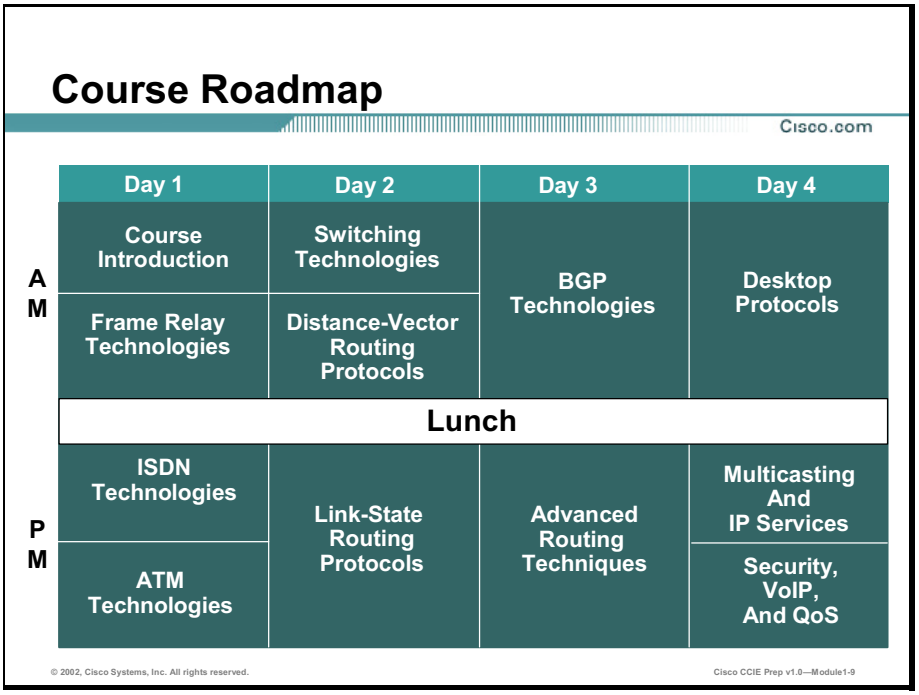

The schedule reflects the recommended structure for this course. This structure allows enough time for the instructor to present the course information and for you to work through the laboratory exercises. The exact timing of the subject materials and labs depends on the pace of your specific class.

Icons and Symbols

This section shows the Cisco icons and symbols used in this course.

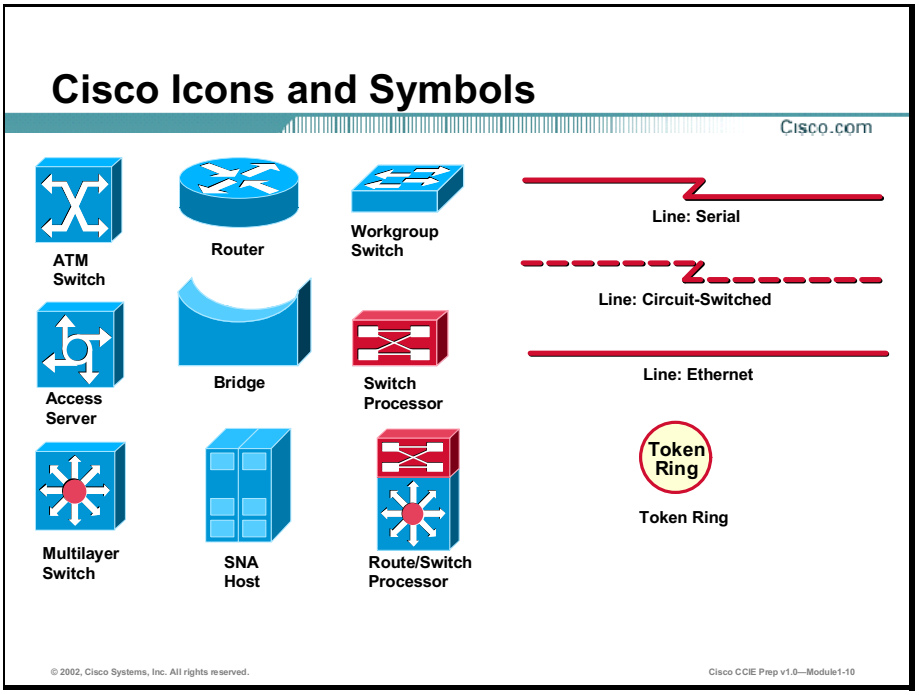

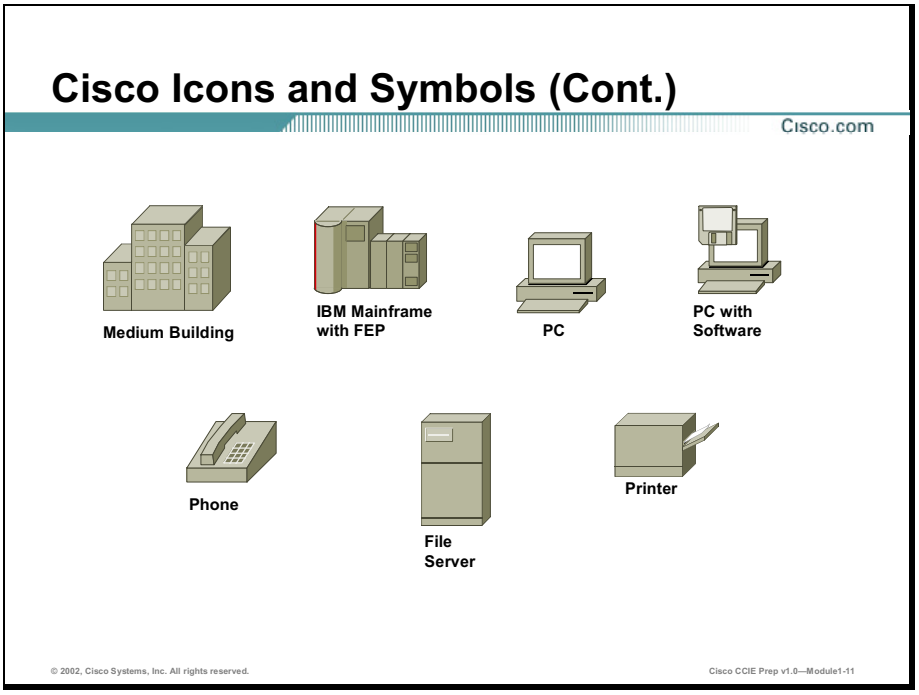

Learner Introductions

This is the point in the course where you introduce yourself.

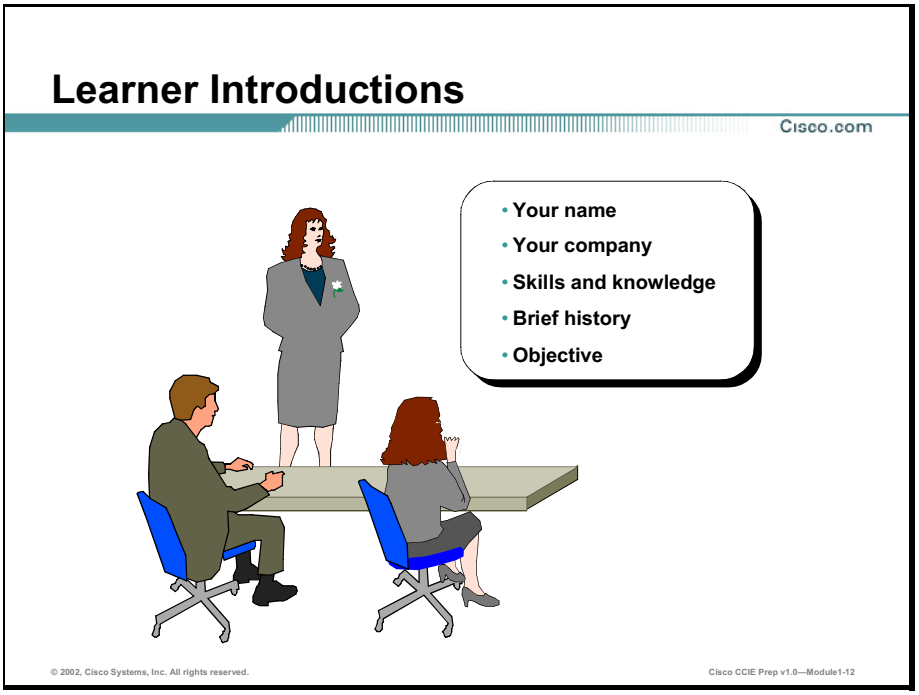

Prepare to share the following information:

- Your name
- Your company
- A profile of your experience
- What is your lab date?
- Are there any subject areas you would like to concentrate on?

Lab Registration

This section covers lab registration.

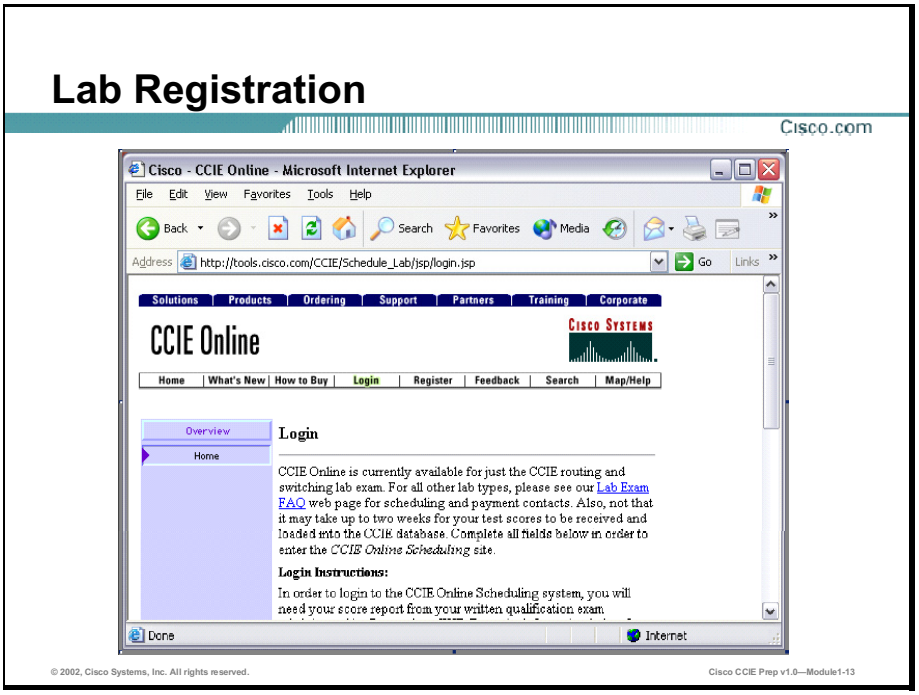

You can now register for the CCIE exam through an on-line system accessible through the Internet. To register for the CCIE Routing and Switching Lab Exam go to:

http://tools.cisco.com/CCIE/Schedule_Lab/jsp/login.jsp

The registration utility will ask you for your candidate ID, which is usually your Social Security Number, the date you would like to take the exam, and the score you received on your written test.

What to Expect the Day of the Lab

This section covers what you can expect on the day of the lab.

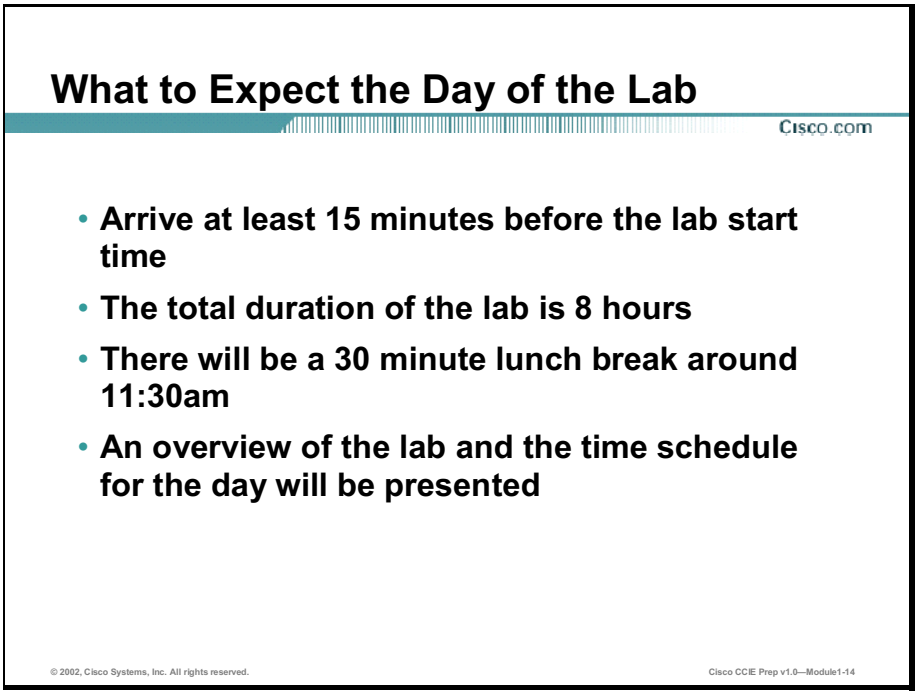

The lab proctor will escort all lab participants into the lab a few minutes before the start time to give everyone an overview of the lab. If you are not present at this time, the proctor may later deny you entry into the lab facilities. The lab personnel will give you an overview of the lab facilities and the time schedule for the day. Cisco should email the lab results to you the following business day.

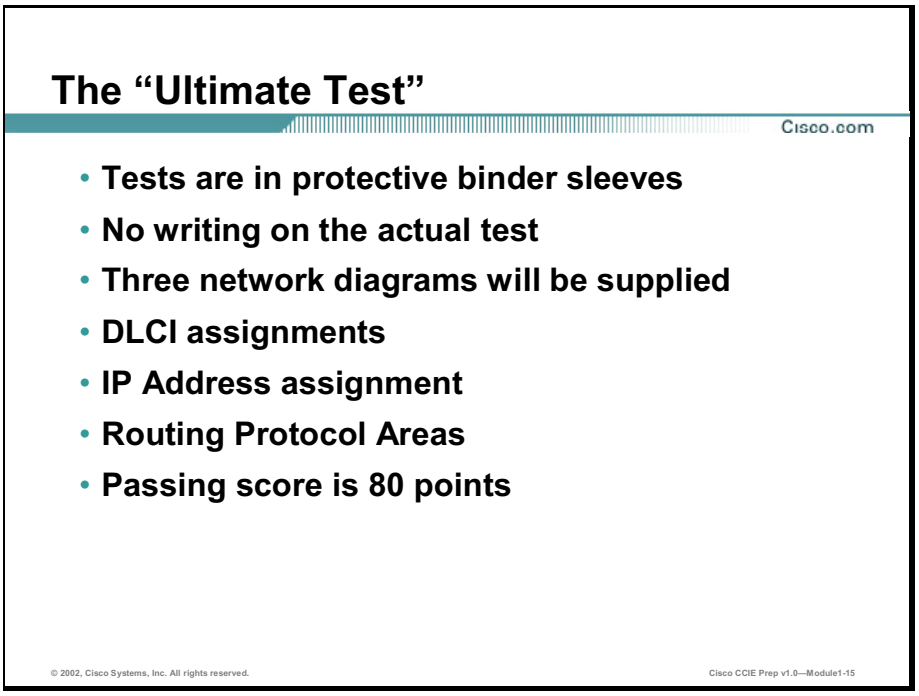

The Lab scenario will be in a loose-leaf notebook. There is no writing on the actual exam. However, scratch paper is available for additional drawings and notes. The proctor will track the number of pieces of scratch paper. Failure to return any pieces of paper will result in automatic failure of the lab exam. The lab includes three network diagrams with the following information: Frame Relay Data-Link Connection Identifier (DLCI) assignments, Internet Protocol (IP) address assignments, and a map of the routing protocols you are to configure. If you need to, you may want to make your own network drawing using the scratch paper provided.

You must score at least 80 points out a possible 100 points to pass the CCIE Lab Exam.

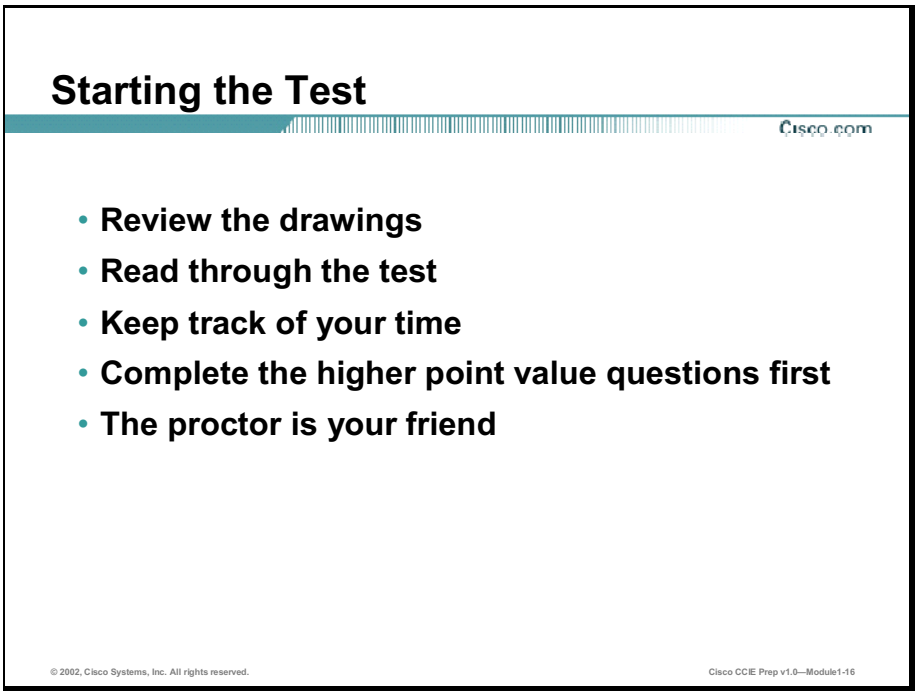

When you start your test, review the drawings and read the test at least once. This will let you know the subject matters covered on your particular exam. It will also help you to determine which questions to answer first. Reading through the test also allows you to configure your devices appropriately based on any future requirements. Requirements at the end of the lab may affect your configuration at the beginning. This is an “issue spotting” test. As you are reading the questions, think about what is involved in configuring the particular scenario and what the implications might be. The lab tests your knowledge of routing protocol interaction.

When configuring a particular question, work methodically. Work based on the Open Systems Interconnection (OSI) model from Layer 1 up. If possible, test each possible answer before proceeding. For example, make sure your Open Shortest Path First (OSPF) neighbors are adjacent and exchanging routing updates before modifying timers or adding authentication. This will allow for easier troubleshooting. Remember you only want to configure a scenario once.

Keep track of your time. If you do not know an answer, move on. Work on the higher point values first. If you are not sure how to answer a question, ask the proctor. Have alternatives available. For instance, tell them “there are two ways to answer this question, method ‘A’ or method ’B’, which would you prefer”. If you suspect a hardware problem, notify the proctor immediately. Be able to substantiate your claim. The proctors are there as a resource, they are not there to help you answer the questions.

Make sure you answer all parts to a question. Allow 30 minutes to one hour before the end of your exam to review your configurations.

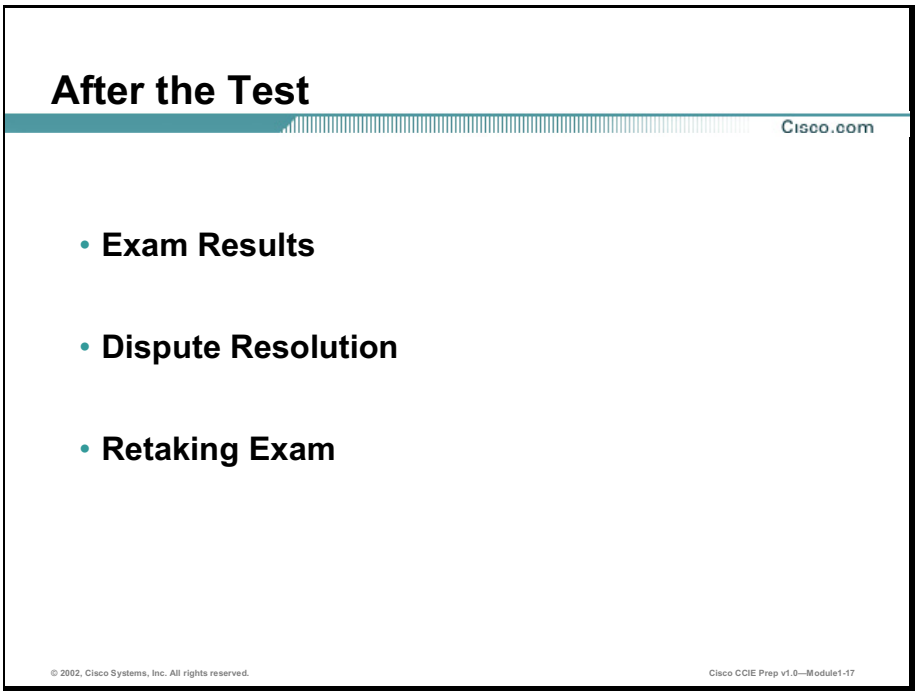

Cisco typically e-mails lab results to you sometime the next business day. If you suspect an error with your test grade, a review procedure is available. There is a fee of $250.00 for this service. If the review results in a passing grade or a difference of 20 points or greater, Cisco will refund the fee. If your overall exam score is less than 20 points, you must wait at least six months before retaking the exam. Otherwise, you may reschedule as soon as you would like.

Again, you must score at least 80 out of a possible 100 points to pass. When you pass your lab you will be in a select group of people in the world who have obtained this prestigious certification. Good luck!

2

Packet Switched Technologies

Overview

Because of their high speed and efficiency, most modern networks employ some type of packet switched technology in their Wide Area Network (WAN) infrastructure. Understanding the concepts and configuration of Frame Relay and Autonomous Transfer Mode (ATM) networks are critical to your success in the Cisco Certified Internetwork Expert (CCIE) lab.

Upon completing this module, you will be able to:

- Describe Frame Relay concepts such as Data-Link Connection Identifiers (DLCI), Inverse Address Resolution Protocol (ARP), and Local Management Interface (LMI)
- Configure Frame Relay on the various interface types: physical, point-to-multipoint subinterfaces, and point-to-point subinterfaces
- Configure remote Layer 3-to-DLCI address mappings, using the following; Inverse ARP, Frame Relay map statements, and specifically assigning a DLCI to a point-to-point subinterface
- Verify your Frame Relay configuration in a layered approach, using the Open Systems Interconnection (OSI) reference model, with various **show** and **debug** commands
- Configure Permanent Virtual Connections (PVCs) on a Cisco router
- Configure ATM quality of service settings
- Troubleshoot ATM configurations on a Cisco router

Outline

The module contains these lessons:

- Frame Relay Configuration

- Troubleshooting Frame Relay
- ATM Configuration and Troubleshooting

Frame Relay Configuration

Overview

Frame Relay can be configured in various topologies and across multiple interface types. This lesson will examine the configuration of Frame Relay on Cisco routers as it relates to the Cisco Certified Internetwork Expert (CCIE) lab.

Importance

Frame Relay is the core Wide Area Network (WAN) technology in the CCIE lab. As a CCIE candidate, you must thoroughly understand the differences between the configuration of Frame Relay on physical interfaces, point-to-multipoint subinterfaces, and point-to-point subinterfaces.

Objectives

Upon completing this lesson, you will be able to:

- Configure Frame Relay on a physical interface
- Configure Frame Relay on a point-to-multipoint subinterface
- Configure Frame Relay on a point-to-point subinterface
- Identify the differences between the **frame-relay map** and **frame-relay interface-dlci** commands

Learner Skills and Knowledge

To fully benefit from this lesson, you must have these prerequisite skills and knowledge:

- Passed the CCIE written qualification exam
- Completed the Building Cisco Remote Access Networks (BCRAN) course or have the equivalent knowledge

Outline

This lesson includes these sections:

- Overview
- Physical Interface Configuration
- Subinterface Configuration
- Summary
- Lesson Assessment (Quiz)

Physical Interface Configuration

Frame Relay can be configured on either the physical interface or a subinterface. This section will discuss the process of configuring Frame Relay on the physical interface. This includes how address mappings are configured on a physical interface and the advantages and disadvantages of using a physical interface for Frame Relay.

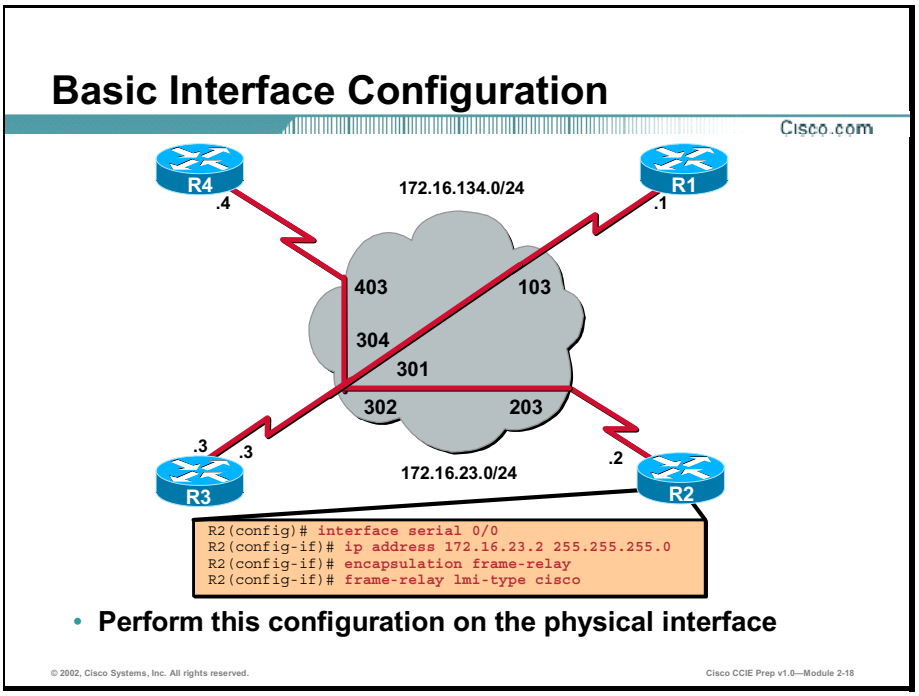

A basic Frame Relay configuration assumes that you want to configure Frame Relay on the physical interface of the router. Perform the following steps to enable Frame Relay on a physical interface:

Table 2-1:

Step	Command	Description
Step 1:	`interface serial 0/0`	Selects the interface and enters interface configuration mode.
Step 2:	`ip address 172.16.23.2 255.255.255.0`	Configures a network-layer address, for example, an IP address.
Step 3:	`encapsulation frame-relay [cisco | ietf]`	Selects the encapsulation type used to encapsulate data traffic end-to-end on the Permanent Virtual Connection (PVC)**,** where **cisco** is the default. Use the default if you are connecting to another Cisco router. Select **ietf** if you are connecting to a non-Cisco router.
Step 4:	`frame-relay lmi-type {ansi | cisco | q933i}`	If using Cisco IOS Release 11.1 or earlier, specify the LMI type used by the Frame Relay switch, where **cisco** is the default. With Internetwork Operating System (IOS) Release 11.2 or later, the LMI type is autosensed and no configuration is needed.

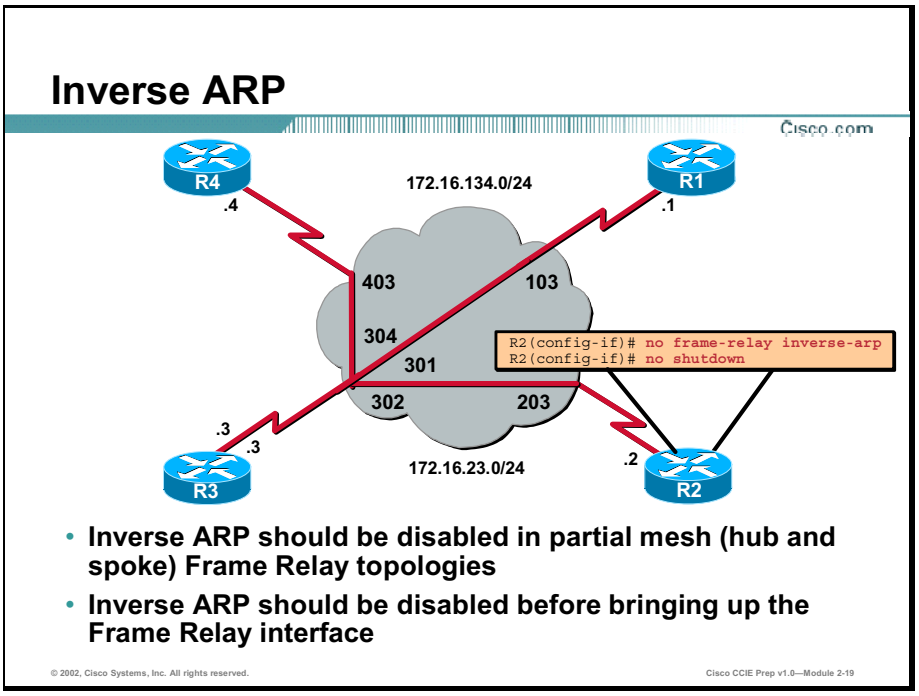

Once the interface is brought up with the **no shutdown** command, the Frame Relay switch will use Local Management Interface (LMI) to communicate the Data-Link Connection Identifier (DLCI) information to the router. Once the DLCIs have attained an active state, meaning that both sides of the connection are up and the Frame Relay switch has the correct Frame Relay route statements, Inverse Address Resolution Protocol (ARP) is performed to map the remote Layer 3 addresses to local DLCIs. Inverse ARP entries are noted in the Frame Relay address mapping table with the keyword dynamic. You can view this table by entering the **show frame-relay map** command.

Inverse ARP works very well in a full mesh Frame Relay topology. However, Inverse ARP has many shortcomings. Inverse ARP will not provide a complete mapping solution in a partial mesh (hub and spoke) topology. Also, Inverse ARP will resolve the Internet Protocol (IP) address of the next-hop router's physical interface even if this IP address is not part of the same IP subnet. This can cause problems in the CCIE Lab and in the real world.

It is recommended that Inverse ARP be disabled on all of your Frame Relay routers in the CCIE Lab. It is also recommended that you disable Inverse ARP before actually bringing up your Frame Relay interfaces with the **no shutdown** command. If Inverse ARP is not disabled before bringing up the interface, you will have to manually clear the Inverse ARP mappings using the **clear frame-relay-inarp** command once your static mappings are in place. Some versions of the Internetwork Operating System (IOS) will actually require a reload to clear the Inverse ARP entries out of the Frame Relay address mapping table, even after entering this command.

Table 2-2:

Step	Command	Description
Step 1:	`no frame-relay inverse-arp`	Disables the sending of Inverse ARP requests on the physical interface.
Step 2:	`no arp frame-relay`	Prevents the router from responding to Inverse ARP requests on an interface.
Step 3:	`no shutdown`	Brings up the physical interface.

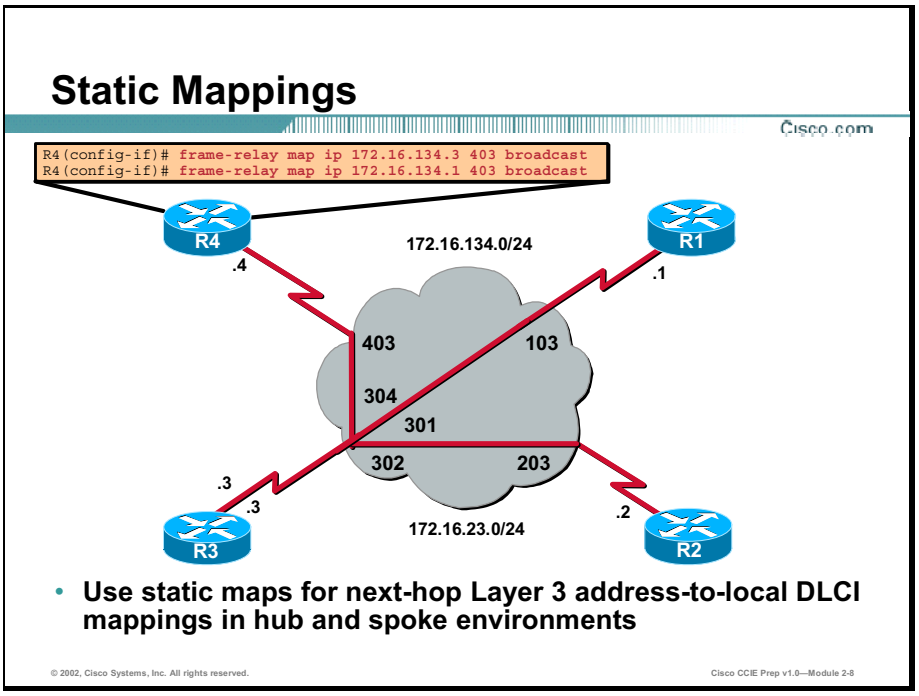

In a hub and spoke topology in which the spoke routers are using physical interfaces, static maps must be used in order for communication between the spokes to occur. A static map links a specified next hop Layer 3 protocol address to a specific DLCI. Static mapping removes the need for Inverse ARP requests. When you supply a static map, Inverse ARP is automatically disabled for the specified protocol on that DLCI.

Table 2-3:

Command	Purpose
frame-relay map protocol protocol-address dlci *[broadcast] [ietf] [cisco]*	Maps a next hop protocol address to a local DLCI. The encapsulation type [cisco | ietf] can be set on a per-PVC basis using the keywords here
protocol	Selects the protocol type. Supported protocols are: appletalk, clns, decent, ip, ipx, xns, and vines
protocol-address	Specifies the protocol address (not specified for bridged or Connectionless Network Service (CLNS) connections)
dlci	Specifies the DLCI number used to connect to the specified protocol address on the interface
broadcast	(Optional) Specifies that broadcasts/multicasts (routing updates) should be forwarded across this PVC
ietf	(Optional) Enables Internet Engineering Task Force (IETF) encapsulation on this PVC
cisco	(Optional) Enables Cisco encapsulation on this PVC

Note You can greatly simplify the configuration for Open Shortest Path First (OSPF) by adding the optional **broadcast** keyword when configuring your Frame Relay map statements. For more information on this see Module 7: Link State Routing Protocols.

Subinterface Configuration

Subinterfaces are used to simplify Frame Relay configurations and resolve reachability issues. Two types of subinterfaces are available, point-to-point and point-to-multipoint. This section will discuss the configuration of both and when to use one type versus the other.

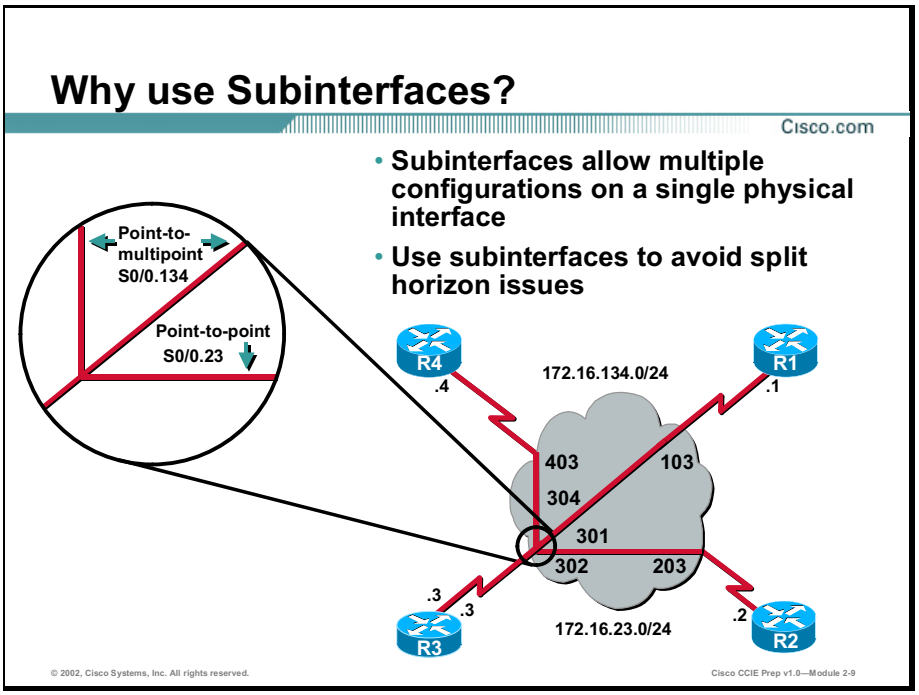

Subinterfaces provide a mechanism for supporting partially meshed Frame Relay networks. Most protocols assume transitivity on a logical network; that is, if R4 can talk to R3, and R3 can talk to R1, then R4 should be able to talk to R1 directly. Transitivity is true on Local Area Networks (LANs), but not on Frame Relay networks, unless R4 is directly connected to R1.

Subinterfaces address the limitations of Frame Relay networks by providing a way to subdivide a partially meshed Frame Relay network into a number of smaller, fully meshed (or point-to-point) subnetworks. Each subnetwork is assigned its own address space and appears to upper layer protocols as if it is reachable through a separate interface.

Configuring Frame Relay subinterfaces ensures that a single physical interface is treated as multiple virtual interfaces, which also allows you to overcome the split horizon rule. Packets received on one virtual interface can be forwarded to another virtual interface, even if they are configured on the same physical interface.

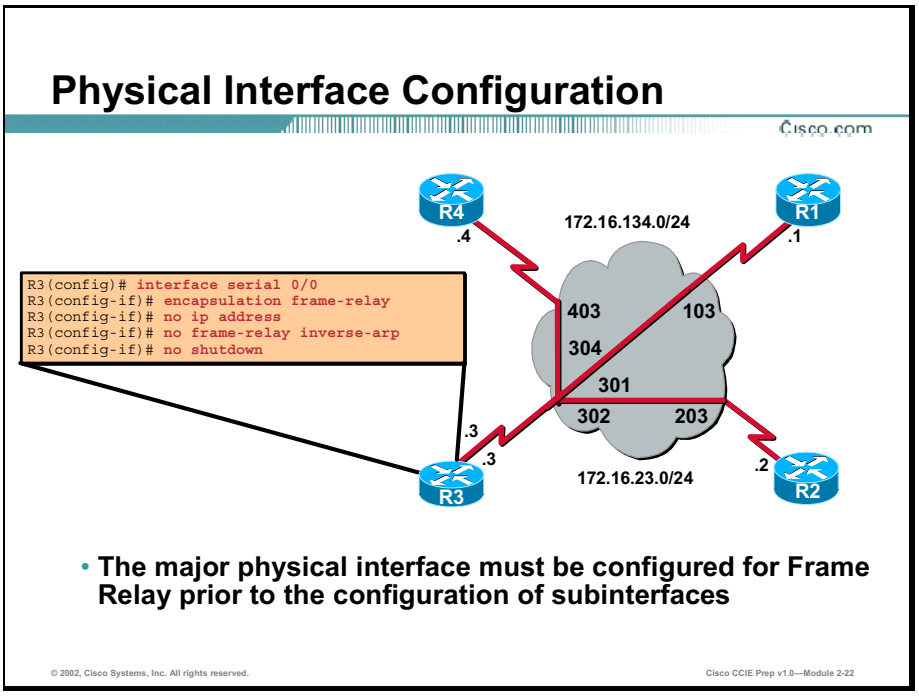

Subinterfaces are typically used at the hub router to resolve reachability issues or simplify Frame Relay configuration. Before we configure subinterfaces, there are some basic configuration steps that must be performed on the physical interface.

Table 2-4:

Step	Command	Description
Step 1:	`interface serial 0/0`	Selects the physical interface that the subinterface will reside under and enters interface configuration mode.
Step 2:	`encapsulation frame-relay [cisco | ietf]`	Enables Frame Relay encapsulation on the physical interface. Frame Relay encapsulation is required for subinterfaces.
Step 3:	`no ip address`	Makes sure there is no Layer 3 address assigned to the physical interface.
Step 4:	`no frame-relay inverse-arp`	Disables Inverse ARP on the physical interface.
Step 5:	`no shutdown`	Brings up the physical interface.

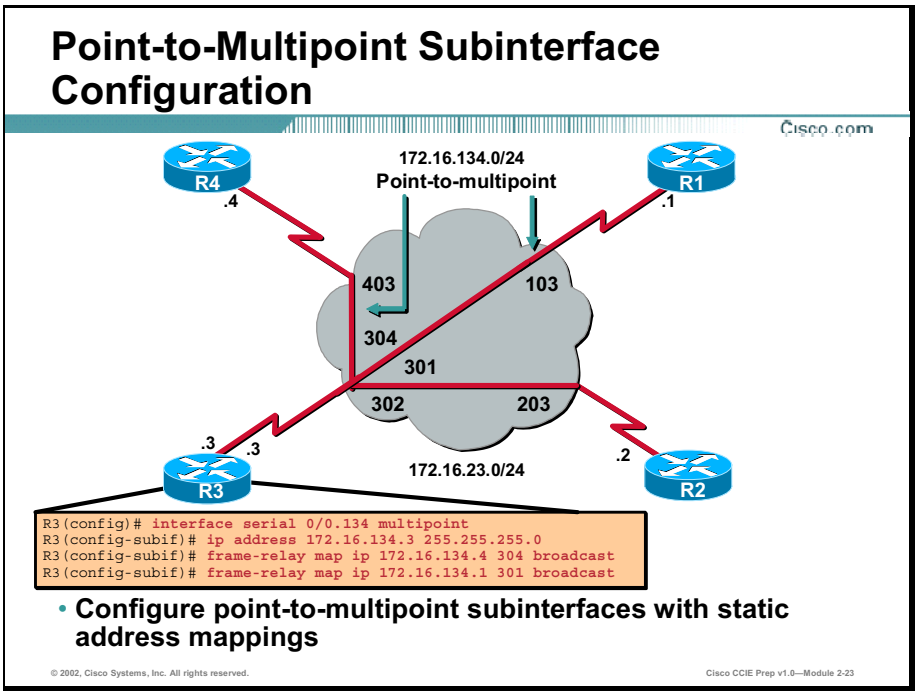

A point-to-multipoint subinterface functions very much like a physical Frame Relay interface. Point-to-multipoint subinterfaces are used to connect spoke routers that reside on the same IP subnet. Frame Relay map statements are used to configure address mappings on point-to-multipoint subinterfaces.

Table 2-5:

Step	Command	Description
Step 1:	`interface serial 0/0.134 multipoint`	Configures a point-to-multipoint subinterface.
Step 2:	`ip address 172.16.134.3 255.255.255.0`	Configures a network-layer address, for example, an IP address.
Step 3:	`frame-relay map ip 172.16.134.4 304 broadcast`	Configures Layer 3-to-DLCI mappings for remote routers.
Step 4:	`frame-relay map ip 172.16.134.1 301 broadcast`	Configures Layer 3-to-DLCI mappings for remote routers.

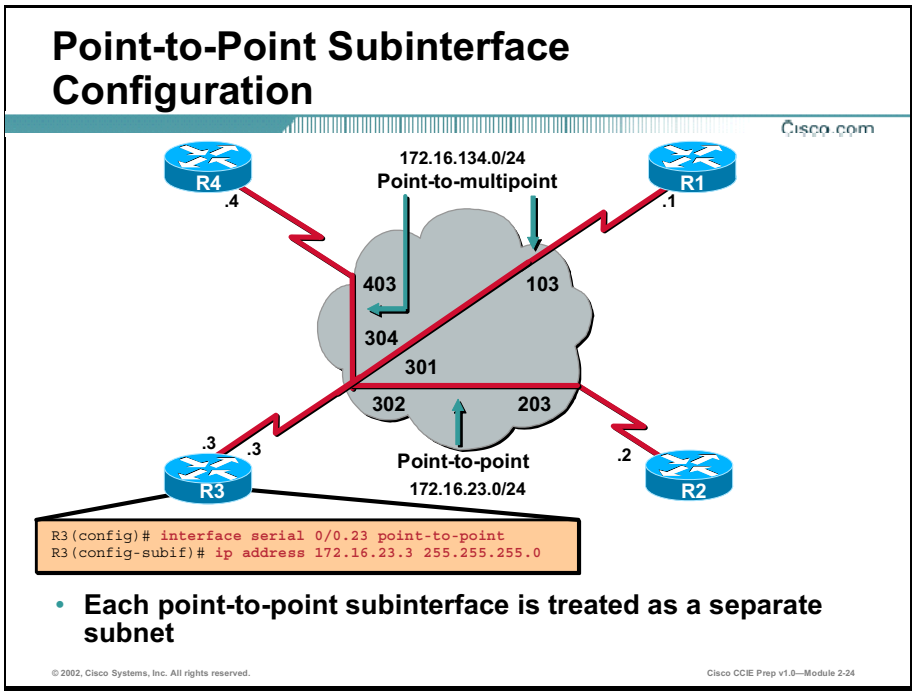

Point-to-point subinterfaces allow a Cisco router to treat each Permanent Virtual Circuit (PVC) as a separate IP subnet. By doing this, the Non-Broadcast Multi-Access (NBMA) characteristics of a Frame Relay network can be avoided. Point-to-point subinterfaces are also used to resolve split horizon issues. By default, split horizon is disabled on all physical interfaces and point-to-multipoint subinterfaces configured for Frame Relay.

This is normally required, as routing updates may need to be received from one spoke router and sent to another on the same interface. This solution has limitations and drawbacks. The limitations are that split horizon can only be disabled for IP and Internetwork Packet Exchange (IPX). One of the drawbacks is that disabling split horizon increases the chances of routing loops in the network when dealing with distance vector routing protocols such as Routing Information Protocol (RIP).

To avoid the problems that come with disabling split horizon, you can use point-to-point subinterfaces. Point-to-point subinterfaces allow routing updates to be received on and sent out the same physical interface, as the router treats each point-to-point subinterface as a separate logical interface. The router actually thinks that the routing updates are coming in on one interface and being sent out a separate interface.

Table 2-6:

Command	Description
`interface serial 0/0.23 point-to-point`	Configures a point-to-point subinterface.
`ip address 172.16.23.3 255.255.255.0`	Configures a network layer address; for example, an IP or IPX address.

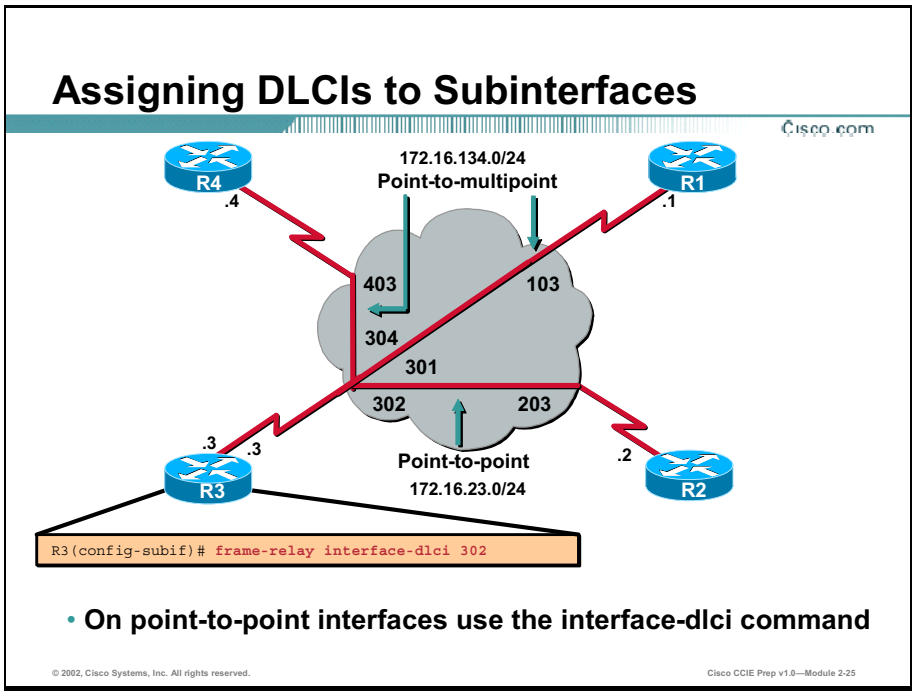

There is no actual remote Layer 3 address-to-DLCI mapping that needs to be configured on a point-to-point subinterface. However, by default, the Frame Relay switch assigns all DLCIs to the physical interface of the Frame Relay Data Terminal Equipment (DTE). Since each point-to-point subinterface is actually a separate PVC, all you need to do is assign the correct DLCIs to the correct subinterfaces. Only one DLCI can be assigned to a particular point-to-point subinterface. The subinterface will then in turn send all Frame Relay traffic down the specified DLCI.

Table 2-7:

Command	Description
`frame-relay interface-dlci` *`dlci`*	Associates the selected point-to-point subinterface with a DLCI

Summary

This section summarizes the key points discussed in this lesson.

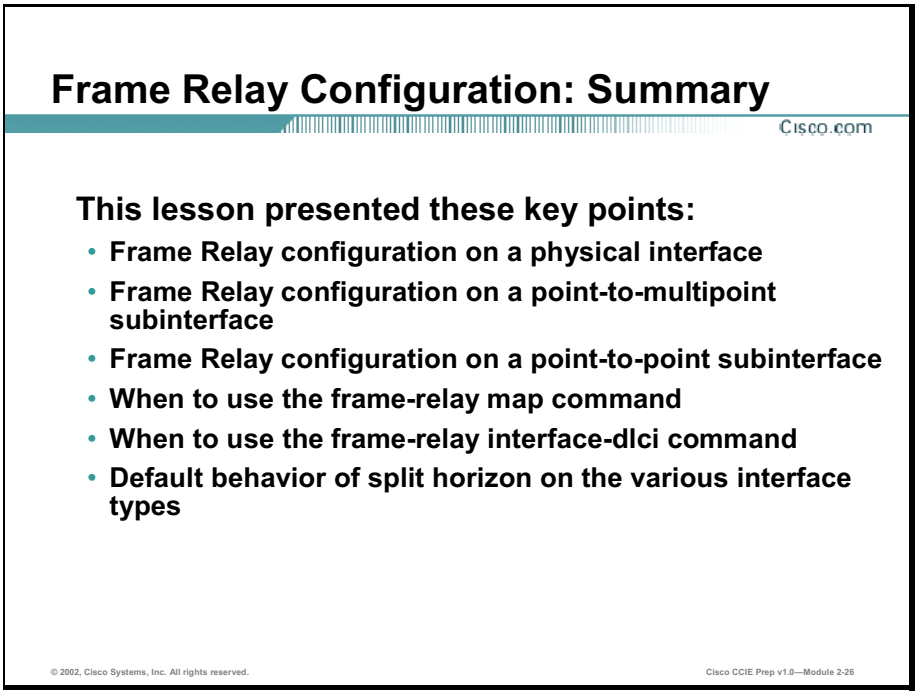

Next Steps

After completing this lesson, go to:

- Troubleshooting Frame Relay

References

For additional information, refer to these resources:

- Building Cisco Remote Access Networks (BCRAN) Module 11
- Configuring Frame Relay - http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121cgcr/wan_c/wcdfrely.htm

Lesson Assessment (Quiz)

Q1) What command is used to clear dynamic Frame Relay mappings learned via Inverse ARP?

Q2) The **frame-relay map** command is used on which of the following interface types?

A) Physical

B) Point-to-multipoint subinterface

C) Point-to-point subinterface

Q3) The **frame-relay interface-dlci** command is used on which of the following interface types?

A) Physical

B) Point-to-multipoint subinterface

C) Point-to-point subinterface

Q4) What does the optional **broadcast** keyword on the **frame-relay map** command do?

Q5) Split horizon for IP is disabled on which of the following interface types by default in a Frame Relay topology?

A) Physical

B) Point-to-multipoint subinterface

C) Point-to-point subinterface

Troubleshooting Frame Relay

Overview

Good troubleshooting skills are necessary for any network engineer. This lesson will teach you about the various **show** and **debug** commands that are available to test and verify Frame Relay configurations.

Importance

Being able to diagnose and quickly troubleshoot problems is a key element of the Cisco Certified Internetwork Expert (CCIE) lab.

Objectives

Upon completing this lesson, you will be able to:

- Verify the status of Layer 1 and Layer 2 using the **show interface** command
- Verify Layer 2 connectivity with the **show cdp neighbors** command
- Use various **show** and **debug** commands to troubleshoot problems in a Frame Relay network
- Verify remote Layer 3 address-to-DLCI mappings with the **debug frame packet** command
- Verify LMI Status messages with the **show frame-relay lmi** and **debug frame-relay lmi** commands

Learner Skills and Knowledge

To fully benefit from this lesson, you must have these prerequisite skills and knowledge:

- Passed the CCIE written qualification exam
- Completed the Building Cisco Remote Access Networks (BCRAN) course or have the equivalent knowledge

Outline

This lesson includes these sections:

- Overview
- Verifying Frame Relay Operation (Layer 1 and 2)
- Verifying Frame Relay Operation (Layer 3)
- Summary
- Lesson Assessment (Quiz)

Verifying Frame Relay Operation (Layer 1 and 2)

This section highlights the various **show** and **debug** commands that are available to verify the operation of Frame Relay at the Physical (Layer 1) and Data-Link Layers (Layer 2) of the Open Systems Interconnection (OSI) model.

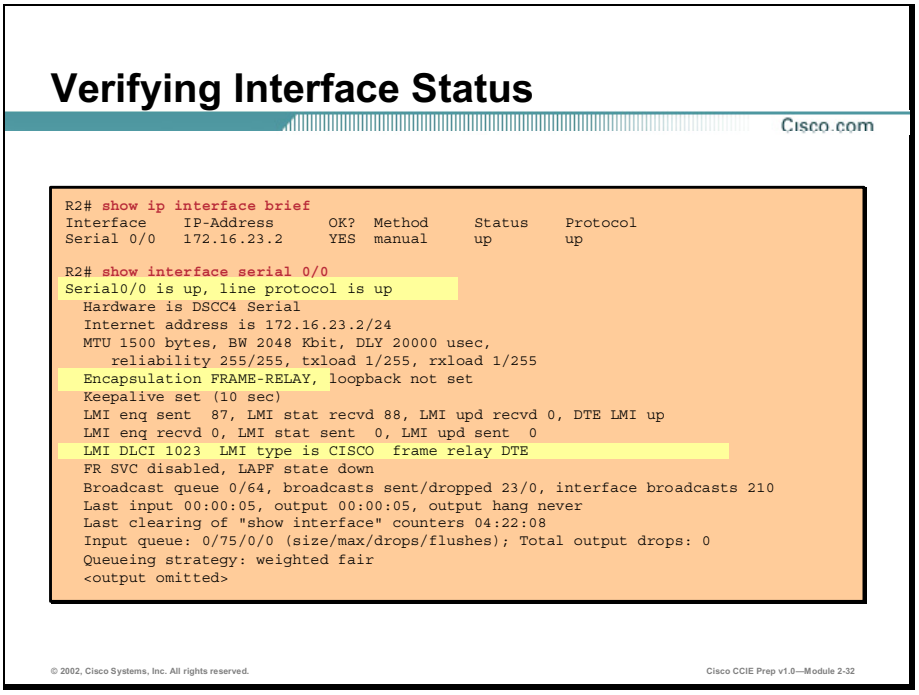

The first thing to check when troubleshooting any type of connectivity problem is the status of the interface. A brief summary of all interfaces (including subinterfaces) can be obtained by issuing the **show ip interface brief** command.

Table 2-8: < show ip interface brief > Command

Command	Description
show ip interface brief	Displays a brief summary of the Layer 1 and Layer 2 status of all IP interfaces on the router

To get detailed information about a specific interface, use the **show interface** command.

Table 2-9: < show interface > Commands

Command	Description
show interface [type] [number]	Displays detailed statistics about the specified interface, including Layer 1 and 2 status, encapsulation type, and LMI information

The items to pay particular attention to in the output of this command are:

- Make sure that the encapsulation is set to Frame Relay.
- If the router is the Data Communication Equipment (DCE) device, make sure you are providing clocking.
- If you are using the Local Management Interface (LMI) type of Cisco, Data-Link Connection Identifier (DLCI) 1023 should be assigned to the physical interface.

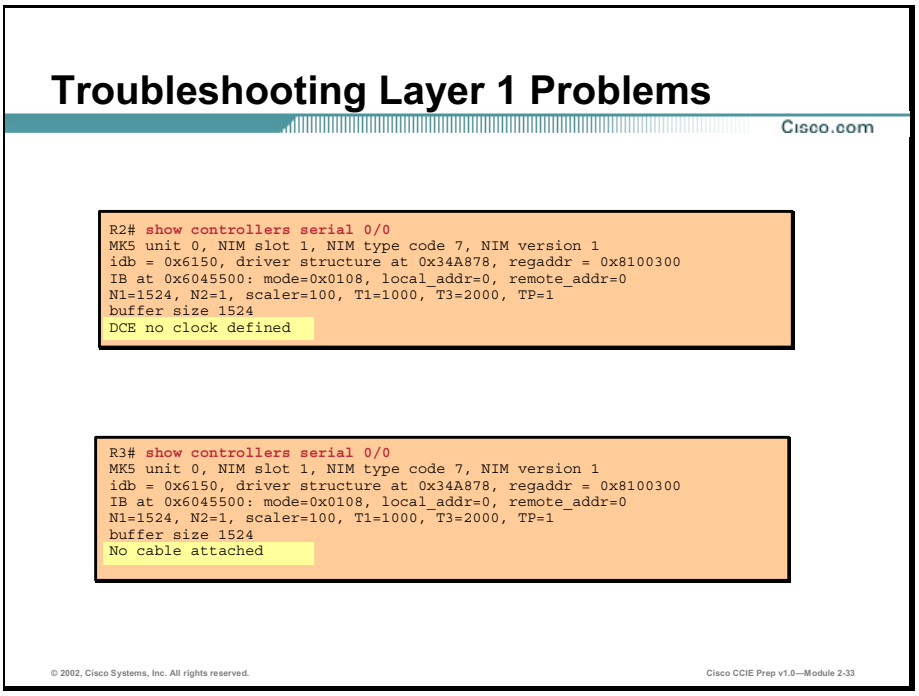

A Layer 1 problem can be determined by the following output from the **show interfaces** command.

```
Serial0/0 is down, line protocol is down
```

Layer 1 problems can usually be tracked down to one of the following: the cable not being plugged in, wrong type of cable, bad cabling, or an interface hardware malfunction.

Most of these problems can be verified using the **show controllers** command.

Note	If you suspect hardware or cabling problems of any kind in the actual CCIE Lab, you should notify your proctor right away.

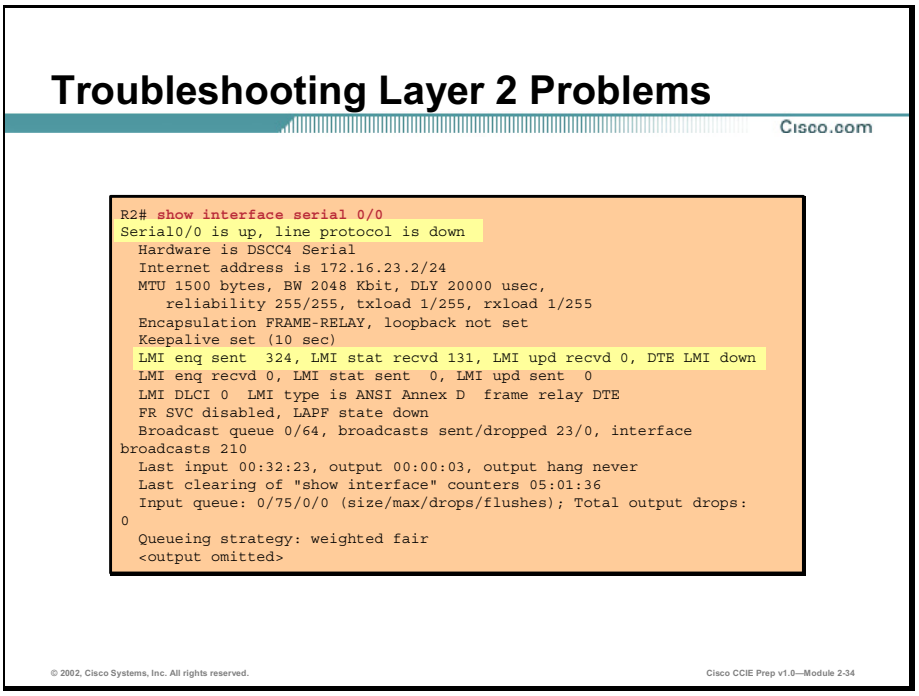

Layer 2 problems are indicated by the following output in the **show interfaces** command:

```
Serial0/0 is up, line protocol is down
```

This usually indicates one of the following: an encapsulation mismatch, no clocking on the link, or when dealing with Frame Relay specifically, it could indicate that LMI is not being received from the Frame Relay switch.

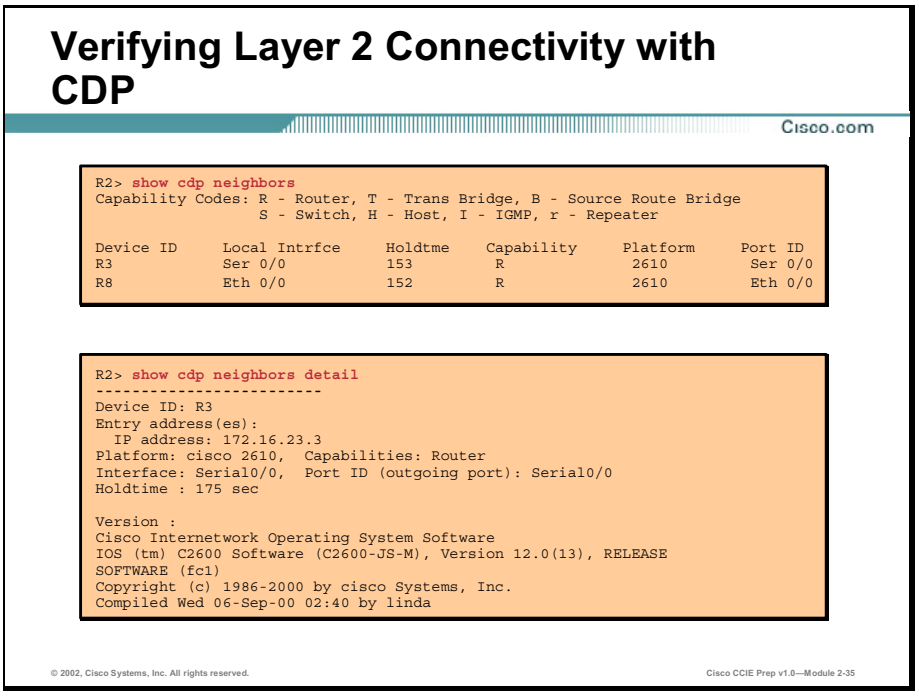

Cisco Discovery Protocol (CDP) is a media and protocol independent protocol that runs on all Cisco-manufactured equipment including routers and switches. Using CDP, you can view information about directly attached devices. CDP runs on all media that supports Subnetwork Access Protocol (SNAP), including Frame Relay. Since CDP works at Layer 2, the following command can be used to verify Layer 2 connectivity to directly connected neighbors.

Table 2-10: <show cdp neighbors> Command

Command	Description
show cdp neighbors	Displays CDP information about directly connected neighbors

CDP can also be used to determine or verify the Layer 3 address of a directly connected neighbor using the following command:

Table 2-11: <show cdp neighbors detail> Command

Command	Description
show cdp neighbors detail	Displays detailed information about a neighbor (or neighbors) including network addresses, enabled protocols, hold time, and software versions

This command can be useful when troubleshooting remote Layer 3 address-to-DLCI mappings.

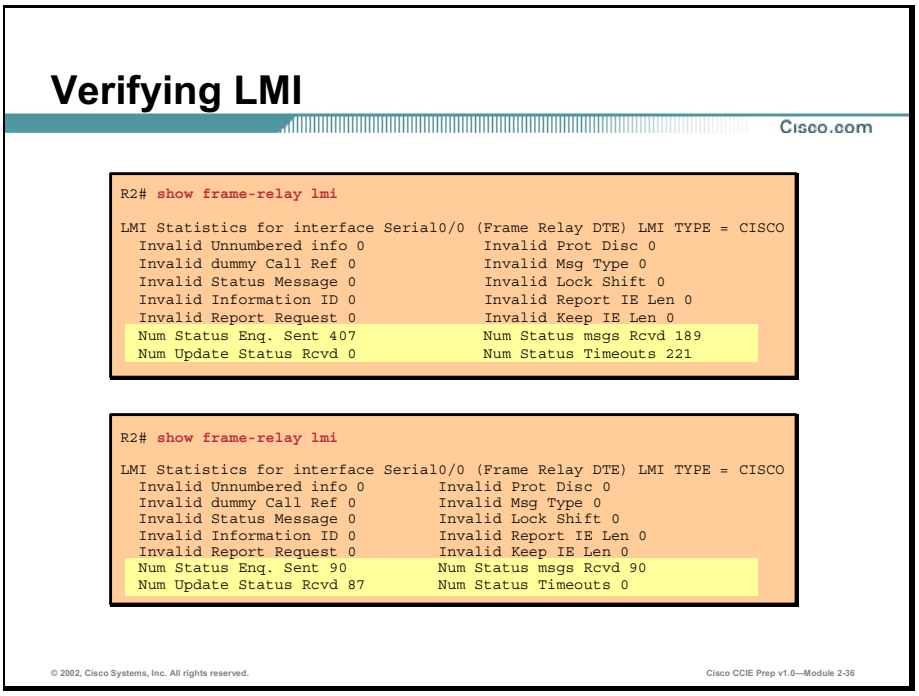

If you are unable to verify Layer 2 connectivity to another router on the same Frame Relay network with the **show cdp neighbors** command, the next step is to verify that the router is communicating with and receiving Local Management Interface (LMI) information from the Frame Relay switch. To do this, enter the following command:

Table 2-12: < show frame-relay lmi [*type number*] > Command

Command	Description
show frame-relay lmi [type number]	Displays LMI statistics
type	Interface type (Optional)
number	Interface number (Optional)

If the router and Frame Relay switch are communicating correctly via LMI, the number of status inquiries sent should match the number of status messages received. Also, the last line in the output of the command (number status timeouts) should be 0 when LMI is functioning correctly.

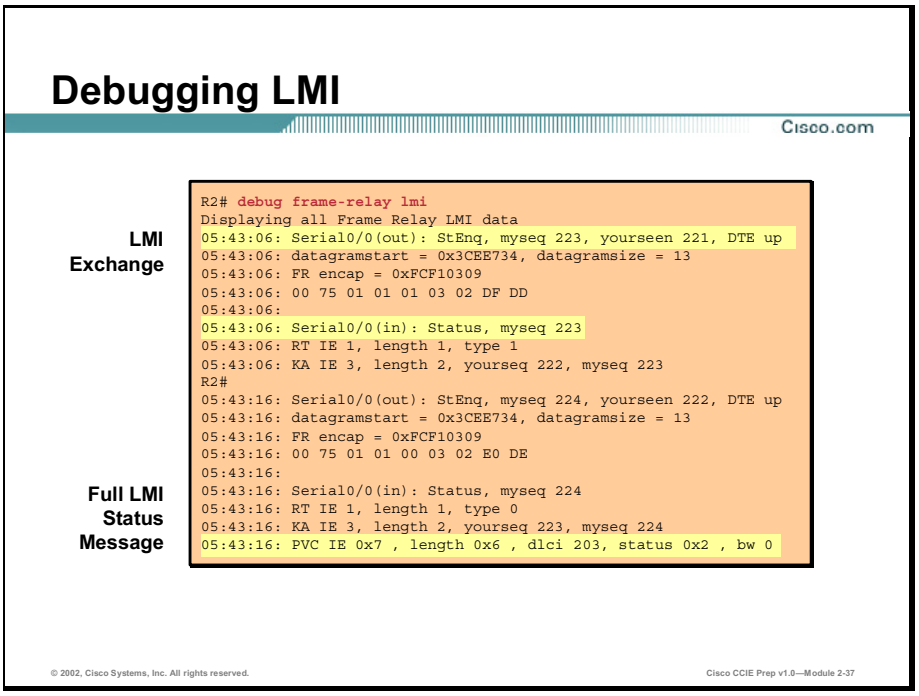

If you notice that Layer 2 is not up on a Frame Relay interface and the number of status enquires sent is not incrementing with the number of status messages received, you should use the following command to determine why:

Table 2-13: < debug frame-relay lmi > Command

Command	Description
debug frame-relay lmi	Displays information on the Local Management Interface (LMI) packets exchanged between the router and the Frame Relay switch

The (out) status field is an LMI status enquiry sent by the router. The (in) status is a reply from the Frame Relay switch.

The type 1 field is a keepalive message sent by the router to the Frame Relay switch approximately every 10 seconds. The purpose of the keepalive message is to verify that the Frame Relay switch is still active. The type 0 field represents Inverse Address Resolution Protocol (ARP) messages exchanged by the routers every 60 seconds.

The dlci 203, status 0x2 field indicates that the status of Data-Link Connection Identifier (DLCI) 203 is active. The possible values of the status field are as follows:

- 0x0 (Inactive) - The switch has this DLCI programmed, but for some reason (such as the other end of the PVC is down) it is not usable.
- 0x2 (Active) - The switch has the DLCI programmed and everything is operational. You can start sending traffic with this DLCI in the header.

- 0x4 (Deleted) – The switch does not have this DLCI programmed for the router. However, it was programmed at some point in the past. This could also be caused by the DLCIs being reversed on the router, or by the Permanent Virtual Circuits (PVC) being deleted by the telco in the Frame Relay cloud.

Note	This debug command is relatively safe to use because full LMI exchanges are only generated every 60 seconds.

Verifying Frame Relay Operation (Layer 3)

This section highlights the various **show** and **debug** commands that are available to verify the operation of Frame Relay at the Network Layer (Layer 3) of the OSI model.

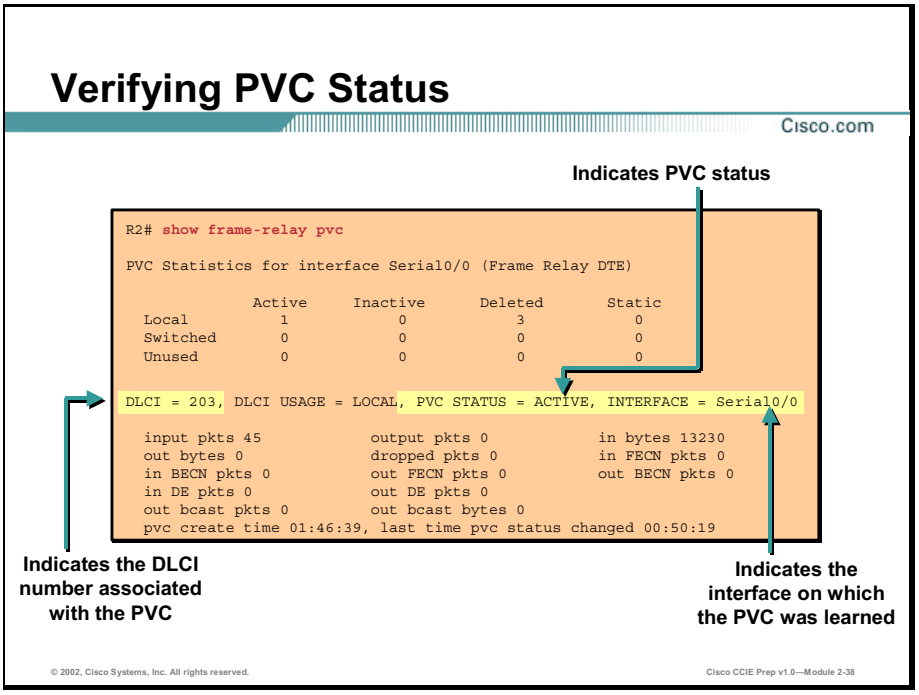

To verify the status of PVCs on the router, enter the following command:

Table 2-14: <show frame-relay pvc [dlci] > Command

Command	Description
`show frame-relay pvc [dlci]`	Displays PVC status

This command also displays parameters dealing with the number of dropped packets, packets marked as Discard Eligible (DE), the number of Backward Explicit Congestion Notifications (BECNs) and Forward Explicit Congestion Notifications (FECNs) received. These items are helpful in verifying the configuration of Frame Relay Traffic Shaping (FRTS). FRTS is covered in Module 12: VoIP, QoS, and Security. Under normal conditions, all PVCs should have a status of "Active".

Table 2-15:

PVC Status	Problem Description
Active	Both sides of the PVC are up and configured properly. The Frame Relay switch also has the correct Frame Relay route statements.
Inactive	This status indicates that the PVC associated with the corresponding DLCI number is being offered by the Frame Relay switch, but not being used by the router.
Deleted	This status indicates that the router has been configured with a DLCI number that is not offered by the Frame Relay switch. As a result, the PVC cannot be created and therefore is "deleted".

If you receive a PVC status of “Inactive” or “Deleted”, double check the DLCI numbering and make certain that the router is configured with the correct DLCIs. DLCI numbers are configured with either the **frame-relay map** command for physical interfaces/multipoint subinterfaces or the **frame-relay interface-dlci** command for point-to-point subinterfaces. A common mistake is the accidental reversal of DLCI numbering. For instance, if the DLCI number that is supposedly assigned to the spoke router shows up as "Inactive" on the hub, there is a good chance that the DLCI numbers are reversed.

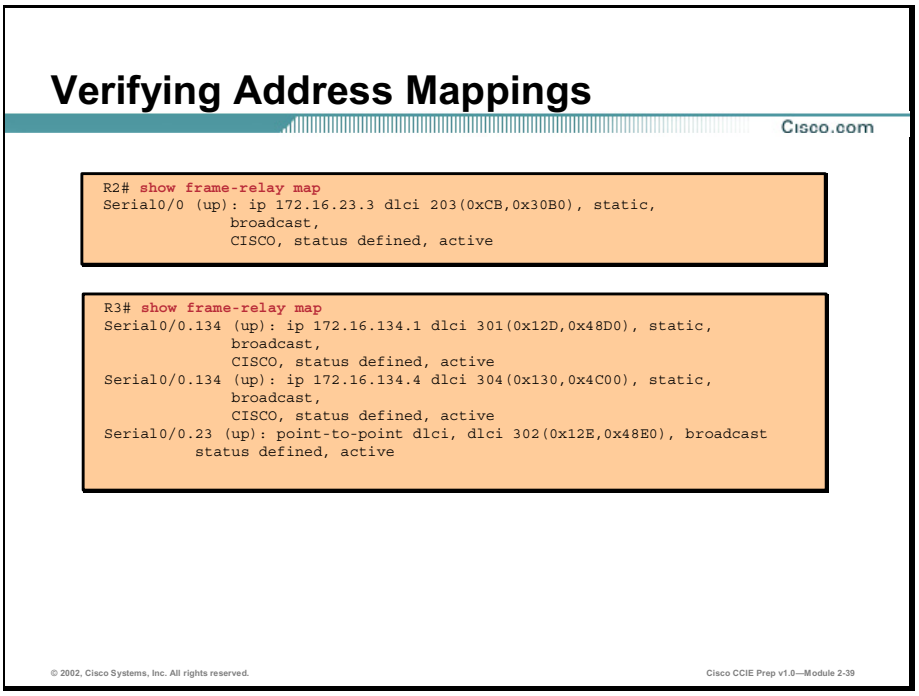

Once you are certain the router is configured with the correct PVCs, you should verify the remote Layer 3 address-to-DLCI mappings. To view the address mapping table on a Cisco router, use the following command:

Table 2-16: < show frame-relay map > Command

Command	Description
`show frame-relay map`	Displays the current address mapping entries

Mappings in this table will either be marked as static or dynamic. Static means that they were statically configured using the **frame-relay map** command. Dynamic indicates that they were learned via Inverse ARP. Point-to-point subinterfaces do not use address mappings and will show up as a point-to-point dlci in the **show frame relay map** command.

If you are not using Inverse ARP and you notice dynamic mappings in the output of the **show frame-relay map** command, you should clear them with the following command:

Table 2-17: < clear frame-relay-inarp > Command

Command	Description
`clear frame-relay-inarp`	Clears dynamically created Frame Relay maps, which are created by the use of Inverse ARP

Note Some versions of the IOS will actually require a reload of the router to clear Inverse ARP entries, even after this command has been entered.

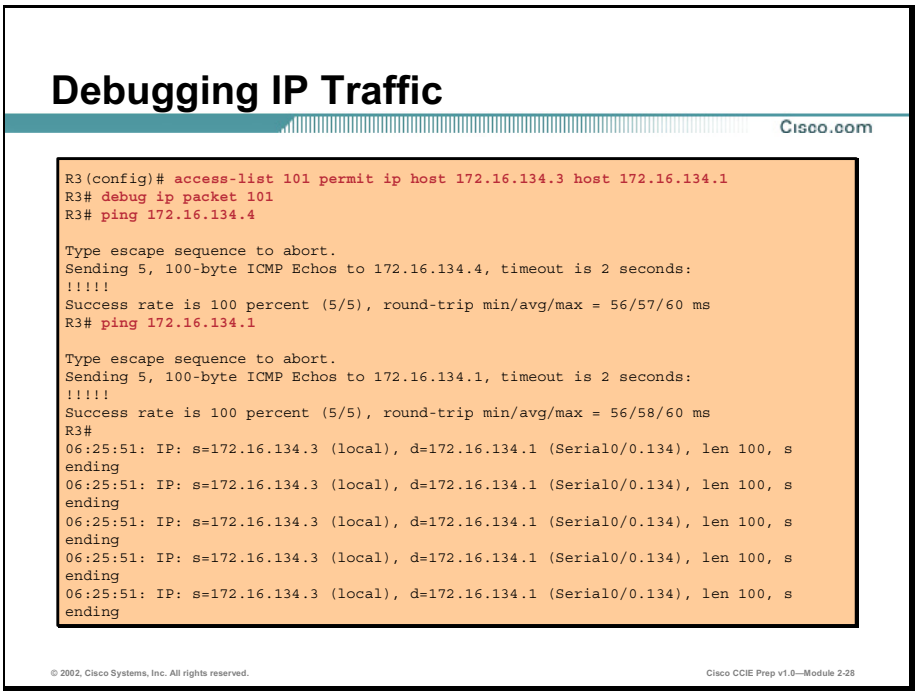

If everything looks good up to this point, you are probably dealing with a Layer 3 or remote Layer 3 address-to-DLCI mapping issue. You can pinpoint this by using the following command:

Table 2-18: < debug ip packet [access-list number] > Command

Command	Description
debug ip packet [*access-list number*]	Displays general IP debugging information
access-list number	IP access list that you can specify. If the datagram is not permitted by that access list, the related debugging output is suppressed.

If you see the following output in the **debug ip packet** command when trying to ping a remote Frame Relay router, you are dealing with a Layer 2 or remote Layer 3 address-to-DLCI mapping issue.

```
IP: s=172.16.1.3 (local), d=172.16.1.1 (Serial0), len 100 sending
IP: s=172.16.1.3 (local), d=172.16.1.1 (Serial0), len 100, encapsulation failed.
```

Note The **debug ip packet** command generates a significant amount of output. Use it only when traffic on the IP network is low, so that other activity on the router is not adversely affected. It is also highly recommend that you use an access list with this command to filter out traffic that you are not interested in debugging.

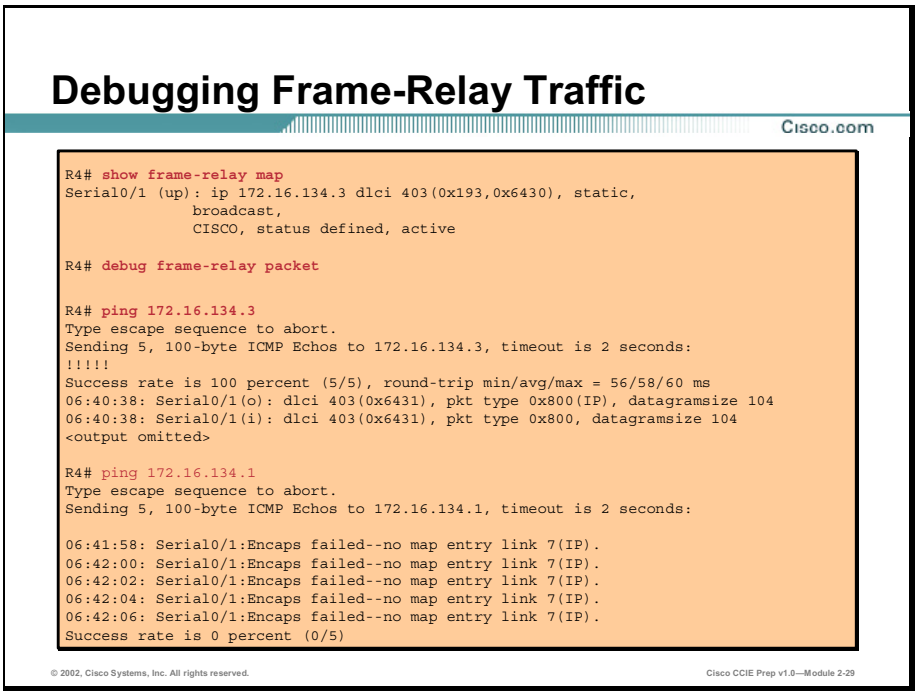

To see what is happening at the packet level of a Frame Relay packet in real-time use the following command:

Table 2-19: < debug frame-relay packet [*interface* [dlci value]] > Command

Command	Description
debug frame-relay packet [*interface* [dlci value]]	Displays information on packets that have been sent and received on a Frame Relay interface

This command allows you to analyze packets that are sent across a Frame Relay interface. Due to the fact that the **debug frame-relay packet** command generates large amounts of output, use it only when traffic on the Frame Relay network is less than 25 packets per second. Additionally, you should use the optional keywords to limit the debugging output to a specific DLCI or interface.

This command is very useful in verifying the configuration of remote Layer 3 address-to-DLCI mappings.

The following line in the output of the command indicates that no address mapping exists for the destination IP address.

Serial0:Encaps failed—no map entry link 7(IP)

Table 2-20:

Field Descriptions - debug frame-relay packet	Description
serial0:	Interface that has been sent to the Frame Relay packet
broadcast = 1	Destination of the packet. Possible values include the following: ■ broadcast = 1—Broadcast address ■ broadcast = 0—Particular destination ■ broadcast search—Searches all Frame Relay map entries for this particular protocol that include the keyword **broadcast**
link 809B	Link type, as documented in the **debug frame-relay** command
addr 172.16.1.1	Destination protocol address for this packet. In this case, it is an IP address.
Serial0(o):	(o) indicated that this is an output event
DLCI 500	Decimal value of the DLCI
type 809B	Packet type, as documented in the **debug frame-relay** command
size 24	Size of this packet (in bytes)

Summary

This section summarizes the key points discussed in this lesson.

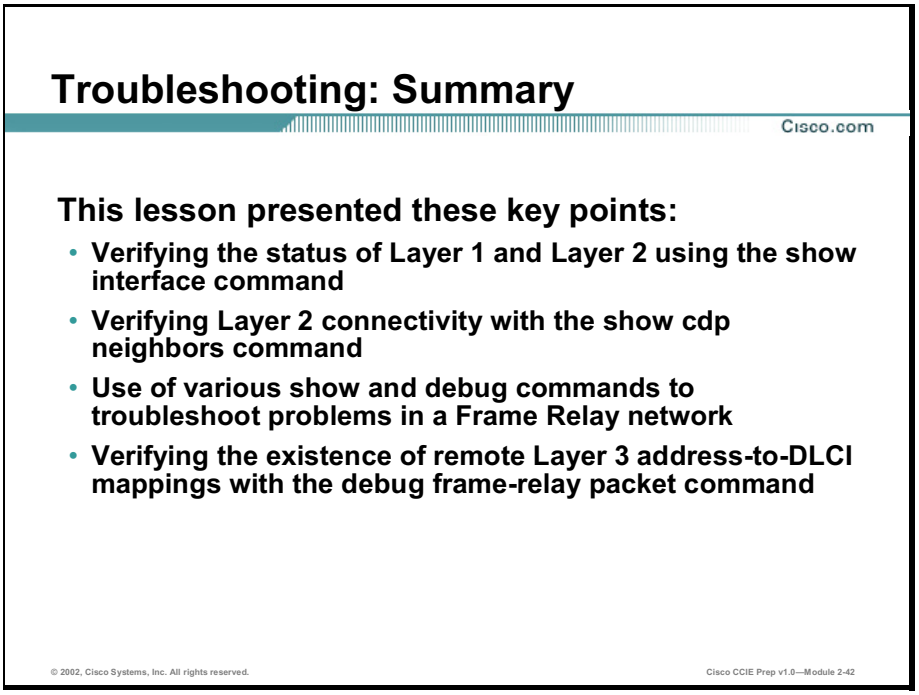

Next Steps

After completing this lesson, go to:

- Integrated Services Digital Network (ISDN) Technologies

References

For additional information, refer to these resources:

- Troubleshooting Frame Relay Connections - http://www.cisco.com/univercd/cc/td/doc/cisintwk/itg_v1/tr1918.htm

Lesson Assessment (Quiz)

Q1) Which of the following indicates a Layer 2 problem?

A) Serial0/0 is up, line protocol is down

B) Serial0/0 is down, line protocol is down

C) Serial0/0 is administratively down, line protocol is down

D) Serial0/0 is up, line protocol is up

Q2) Which of the following indicates a Layer 1 problem?

A) Serial0/0 is up, line protocol is down

B) Serial0/0 is down, line protocol is down

C) Serial0/0 is administratively down, line protocol is down

D) Serial0/0 is up, line protocol is up

Q3) What command is used to verify Layer 2 connectivity to a directly connected neighbor?

Q4) Which debug command is used to verify the existence of a Frame Relay map statement when sending pings to a particular next-hop Layer 3 address?

ATM Configuration and Troubleshooting

Overview

This lesson provides an overview of Autonomous Transfer Mode (ATM) the methods used to configure an ATM Permanent Virtual Connection (PVC). You can use this information to configure ATM connections in the Cisco Certified Internetwork Expert (CCIE) lab and allow routing protocol data to transmit across the link. In addition, this lesson describes the configuration of ATM traffic shaping by using the predefined service classes.

Importance

ATM is a Wide Area Network (WAN) technology tested in the CCIE lab. As a CCIE candidate, you should be able to demonstrate the ability to configure ATM PVCs, configure routing protocols over ATM, and configure QoS on an ATM virtual circuit.

Objectives

Upon completing this lesson, you will be able to:

- Describe what ATM is and why it is a reliable solution for data transmission across the WAN
- Configure ATM PVCs and PVC auto-discovery
- Allow routing traffic to pass over ATM circuits
- Configure ATM Quality of Service (QoS)

Learner Skills and Knowledge

To fully benefit from this lesson, you must have these prerequisite skills and knowledge:

- Passed the CCIE written qualification exam
- Have a firm understanding of WAN technologies, such as Frame Relay

Outline

This lesson includes these sections:

- Overview
- ATM Fundamentals
- ATM Virtual Connections
- Routing over ATM
- Configuring the AAL and Encapsulation Type
- Configuring PVC Traffic Parameters
- Troubleshooting ATM
- Summary
- Assessment (Quiz): ATM Configuration

ATM Fundamentals

This section provides an overview of Autonomous Transfer Mode (ATM) concepts and components.

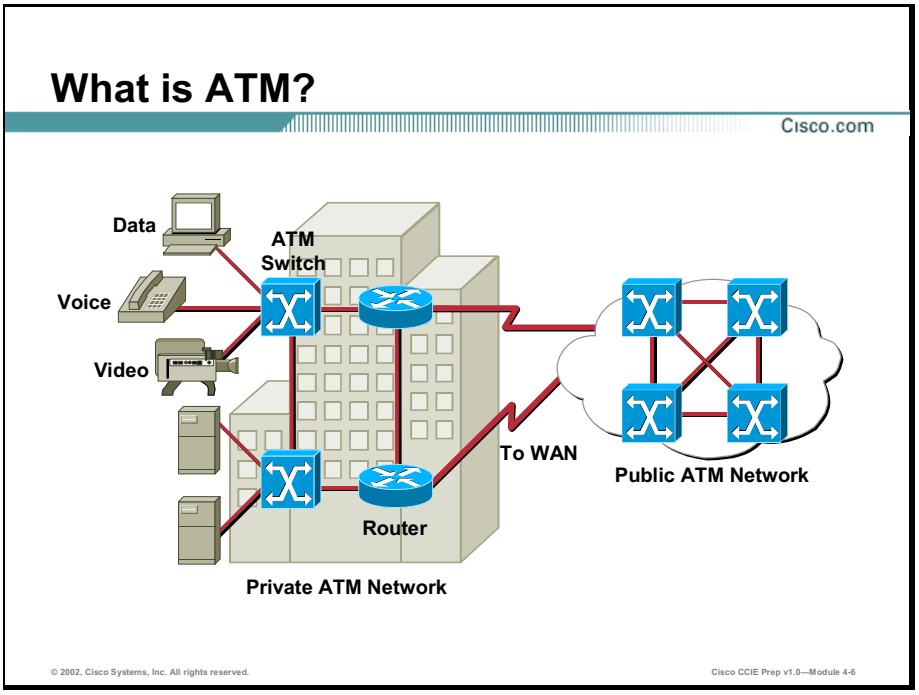

Asynchronous Transfer Mode (ATM) is an International Telecommunication Union Telecommunication Standardization Sector (ITU-T) standard for cell relay wherein routers convey information for multiple service types, such as voice, video, or data, in small, fixed-size cells. ATM networks are connection oriented.

ATM is a cell-switching and multiplexing technology that combines the benefits of circuit switching (guaranteed capacity and constant transmission delay) with those of packet switching (flexibility and efficiency for intermittent traffic). It provides scalable bandwidth from a few Megabits per second (Mbps) to many Gigabits per second (Gbps). Because of its asynchronous nature, ATM is more efficient than synchronous technologies, such as Time-Division Multiplexing (TDM).

With TDM, the router assigns each user to a time slot, and no other station can send in that time slot. If a station has large amounts of data to send, it can send only when its time slot comes up, even if all other time slots are empty. If, however, a station has nothing to transmit when its time slot comes up, the router sends an empty time slot, wasting network bandwidth. Because ATM is asynchronous, time slots are available on demand with information identifying the source of the transmission contained in the header of each ATM cell.

ATM Virtual Connections

This section discusses ATM virtual connections and their importance to network communication.

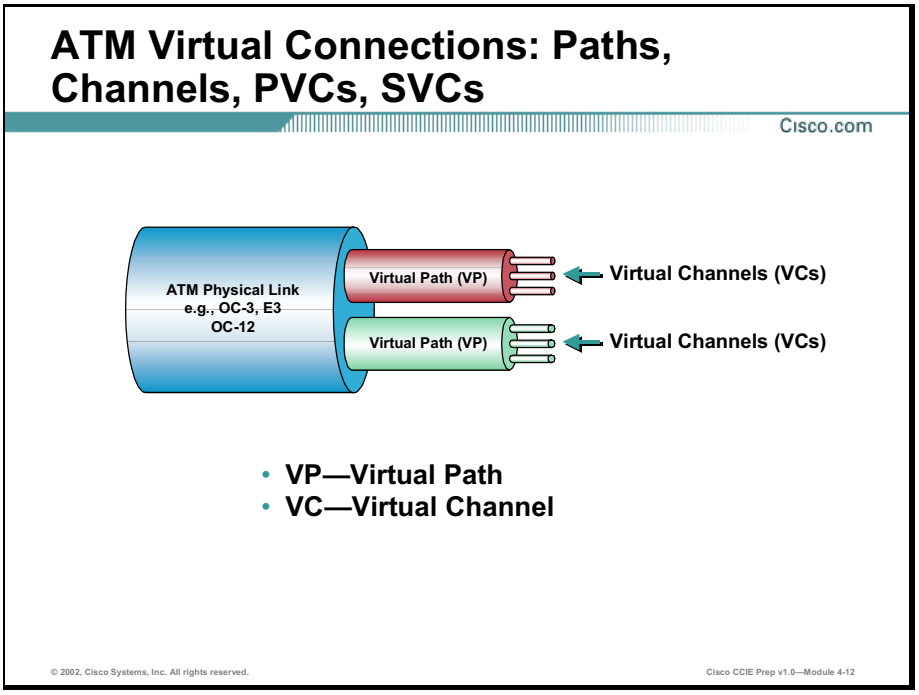

Three types of ATM services exist: Permanent Virtual Circuits (PVCs), Switched Virtual Circuits (SVCs), and connectionless service (which is similar to SMDS).

A PVC allows direct connectivity between sites. In this way, a PVC is similar to a leased line. Among its advantages, a PVC guarantees availability of a connection and does not require call setup procedures between switches. Disadvantages of PVCs include static connectivity and manual setup.

An SVC is created and released dynamically and remains in use only as long as the router is sending data. In this sense, it is similar to a telephone call. Dynamic call control requires a signaling protocol between the ATM endpoint and the ATM switch. The advantages of SVCs include connection flexibility and call setup handled automatically by a networking device. Disadvantages include the extra time and overhead required to set up the connection.

ATM networks are fundamentally connection oriented, which means that the router must set up a Virtual Channel (VC) across the ATM network prior to any data transfer. (A virtual channel is roughly equivalent to a virtual circuit.)

Two types of ATM connections exist: virtual paths, identified by Virtual Path Identifiers (VPIs), and virtual channels, identified by the combination of a VPI and a Virtual Channel Identifier (VCI).

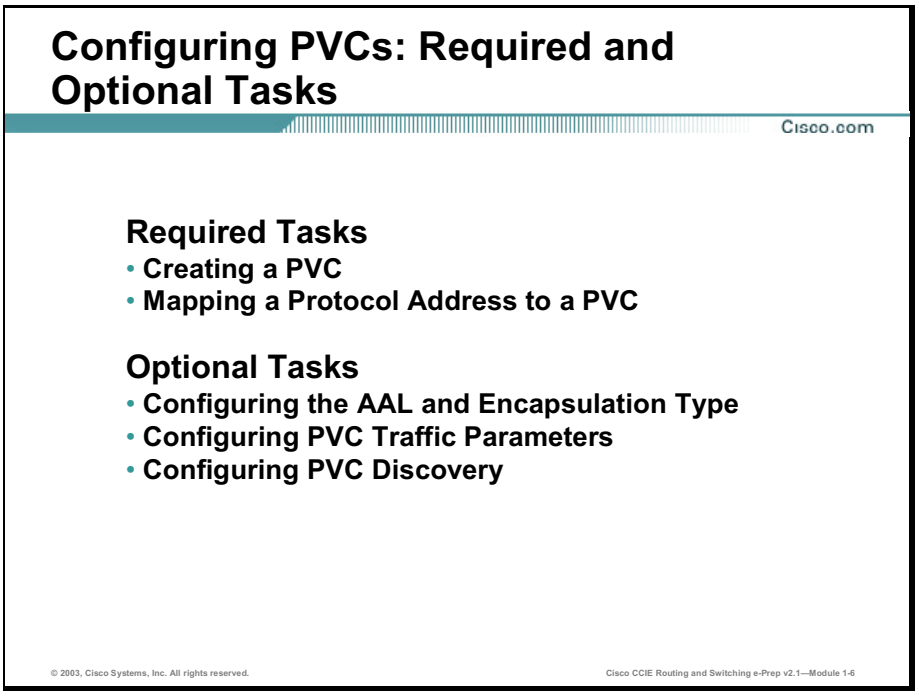

To configure a PVC, perform the following tasks. The first two tasks are required; the other tasks are optional:

- Creating a PVC ***(Required)***
- Mapping a Protocol Address to a PVC ***(Required)***
- Configuring the ATM Adaptation Layer (AAL) and Encapsulation Type ***(Optional)***
- Configuring PVC Traffic Parameters ***(Optional)***
- Configuring PVC Discovery ***(Optional)***

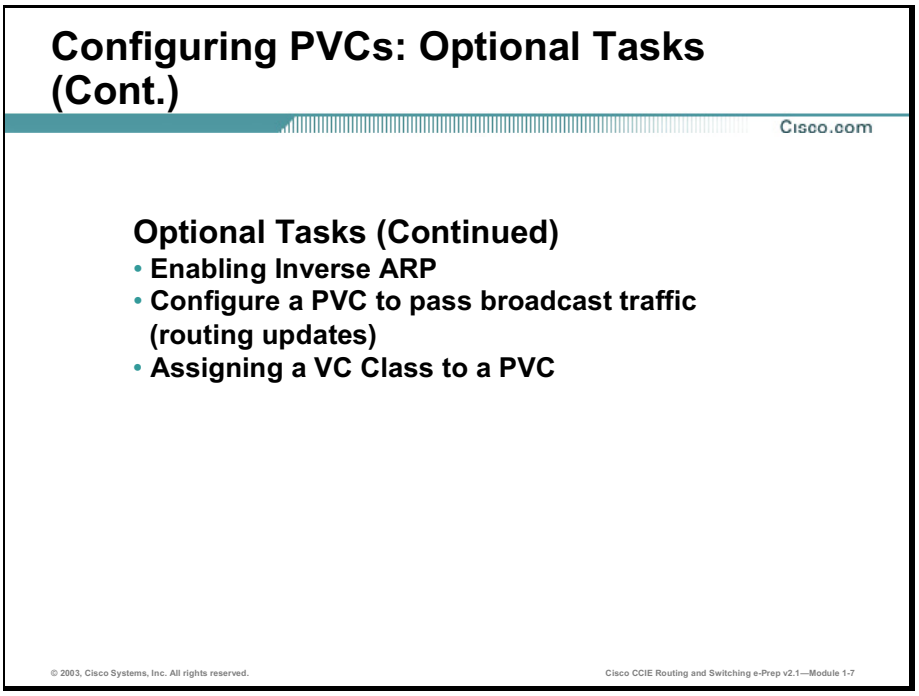

Optional tasks (continued):

- Enabling Inverse Address Resolution Protocol (ARP) ***(Optional)*****:** Allows for dynamic protocol mapping between an ATM PVC and a network address. The router learns the network address dynamically because of the exchange of ATM Inverse ARP packets.

- Configuring a PVC to pass broadcast traffic (routing updates) ***(Optional)*****:** Allows you to send duplicate broadcast packets for all protocols configured on a PVC, using the broadcast keyword in interface-ATM-VC configuration mode.

- Assigning a VC Class to a PVC ***(Optional)*****:** By creating a VC class, you can preconfigure a set of default parameters that you may apply to a PVC.

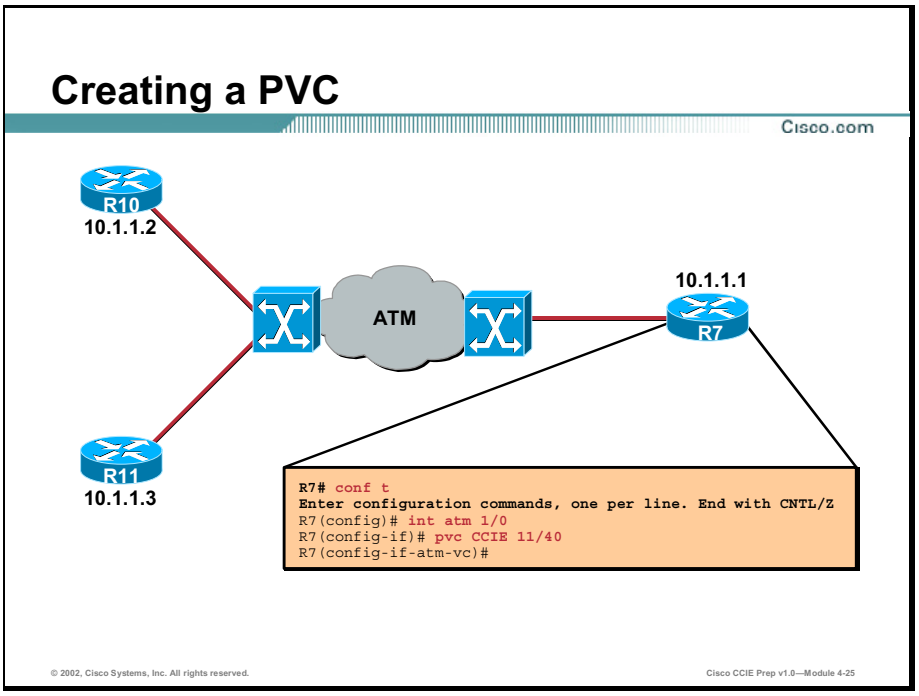

To create a PVC on the ATM interface, use the following command:

pvc [name] vpi/vci [**ilmi** | **qsaal** | **smds**]

As soon as you enter this command, you will be in ATM VC interface configuration mode, where you can specify the parameters of the PVC.

Take a closer look at the general form of the **pvc** command:

Step 1 The mandatory elements of the command are the keyword **pvc** and the *vpi/vci* values.

Step 2 The optional [*name*] parameter allows you to give the PVC a name, which can be used as a handle for further reference or further configuration: once you specify a name for a PVC, you can reenter the interface-ATM-VC configuration mode by simply entering **pvc** *name*.

Step 3 The key-word **ilmi** creates the Integrated Local Management Interface (ILMI) PVC with the Virtual Channel Identifier (VCI) value of 16 needed if you want to configure SVCs. The ILMI PVC is used for ILMI (Integrated Local Management Interface) messages between the end user and the nearest ATM switch.

Step 4 The key-word **qsaal** creates the signaling PVC with the VCI value of 5 needed in an SVC environment.

Step 5 The key-word **smds** allows the intended PVC to handle Switched Multimegabit Data Services (SMDS) over your ATM network.

Note Whenever a PVC is created, the router automatically assigns it a number, designated as VCD = Virtual Circuit Descriptor. In the output from **show** commands, if you name the PVC, you will see its name; if you do not name it, you will see its VCD.

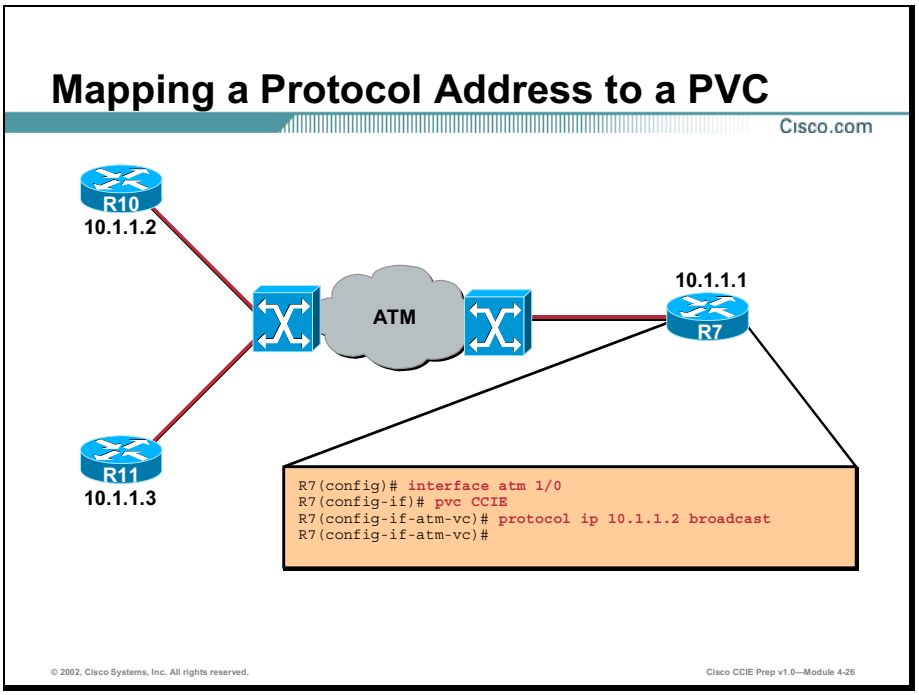

The second required task for the PVC setup is mapping a Layer 3 protocol address to the PVC. You are actually using a static scheme that identifies the network address of the destination host(s).

In (config-if-atm-vc) mode, you will use the following command:

protocol *protocol protocol_address* [[**no**] **broadcast**]

Specify the **broadcast** keyword if you are planning to run any type of routing protocol across this PVC.

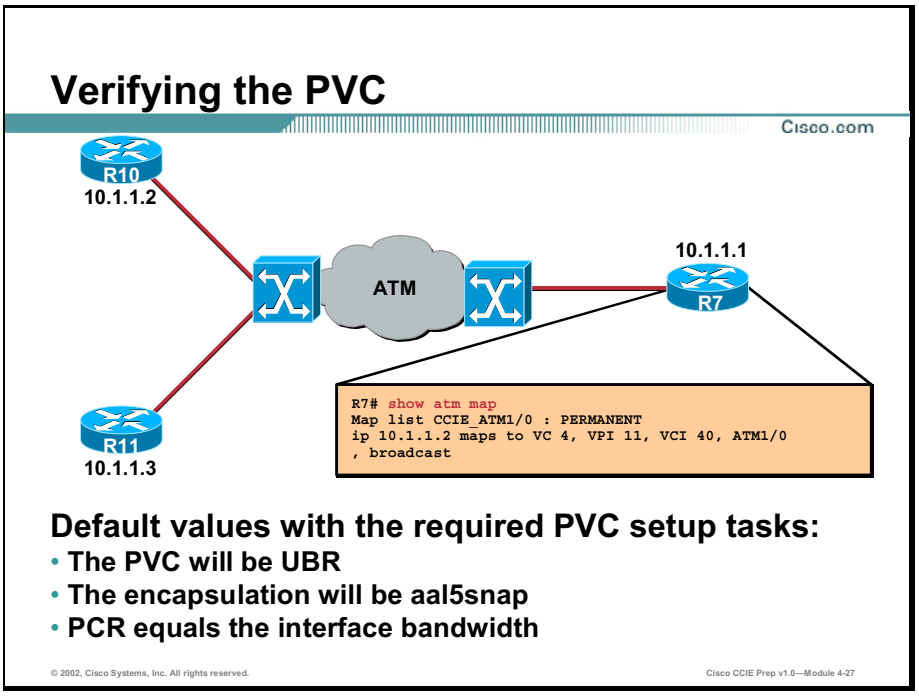

If you stop after configuring the two required tasks described above and check your PVC, you will have created an Unspecified Bit Rate (UBR) PVC whose peak is the physical capacity of the interface, with the Subnet Access Protocol (SNAP) encapsulation type. UBR, the peak physical capacity, and aal5snap are defaults. To illustrate this, use the following **show** commands:

show atm vc [**interface** *slot/module/port*]

show atm pvc [**interface** *slot/module/port* [*vpi/vci* | *vci* | *name*]]

A shorter form of the second command is:

show atm pvc *name*

To verify the second task, mapping a protocol address to a PVC, you will use the command:

show atm map

The output will look like this:

```
R7# show atm map
Map list CCIE_ATM1/0 : PERMANENT
ip 10.1.1.2 maps to VC 4, VPI 11, VCI 40, ATM1/0
, broadcast
```

The router refers to the PVC as VC 4, when it says “ip 10.1.1.2 maps to VC 4”; it has automatically assigned the PVC a Virtual Circuit Descriptor (VCD) number of 4.

Other information you acquire concerns the Virtual Path Identifier/Virtual Channel Identifier (VPI/VCI) values, the interface, and the “broadcast” feature. It also tells you this is a permanent mapping.

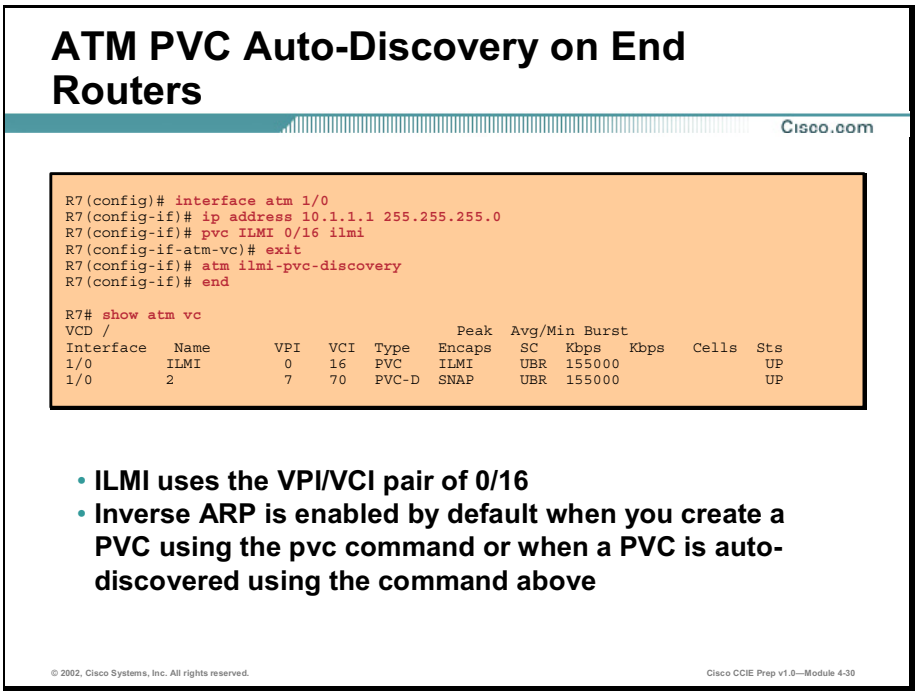

You can use Integrated Local Management Interface (ILMI) to discover and configure the PVCs. ILMI uses what it gets from the adjacent switch. The router discovers the PVCs configured on the switch and configures the PVCs and their traffic parameters on the ATM main interface or subinterface that you specify.

In order to use this feature, first configure the PVC that ILMI will use, with the following commands in interface configuration mode:

pvc [*name*] **0/16 ilmi**

atm ilmi-pvc-discovery [subinterface]

Inverse ARP is enabled by default when you create a PVC using the **pvc** command or when a router auto-discovers a PVC by using the previous syntax. As a result, a configured protocol mapping between an ATM PVC and the router learns a network address dynamically because of the exchange of ATM Inverse ARP packets.

In this example, the service provider has reserved the VPI/VCI values of 7/70 for your PVC.

The router will automatically discover these values using ILMI. Examine the output shown below:

```
R7(config)# interface atm 1/0
R7(config-if)# ip address 10.1.1.1 255.255.255.0
R7(config-if)# pvc ILMI 0/16 ilmi
R7(config-if-atm-vc)# exit
R7(config-if)# atm ilmi-pvc-discovery
R7(config-if)# end
```

```
R7# show atm vc
VCD /                                        Peak  Avg/Min Burst
Interface  Name  VPI  VCI Type  Encaps  SC  Kbps  Kbps  Cells  Sts
1/0        ILMI    0   16  PVC    ILMI UBR  155000              UP
1/0        2       7   70  PVC-D  SNAP UBR  155000              UP

R7# ping 10.1.1.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 4/4/4 ms
```

Routing Over ATM

ATM is a non-broadcast multi-access topology. Therefore, many of the same issues involved with routing over Frame Relay also apply here.

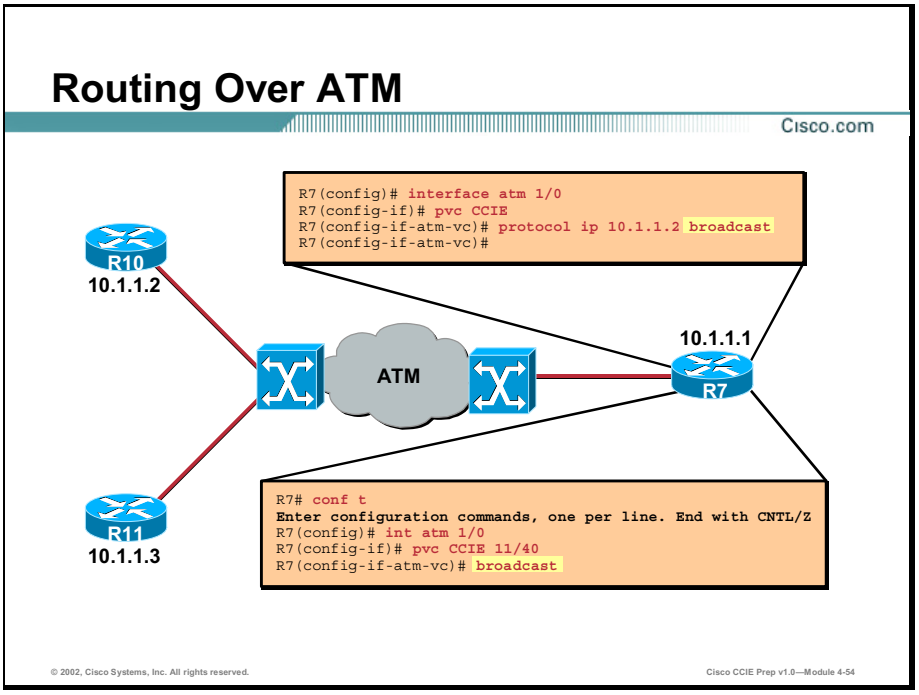

In order for a PVC or SVC to pass routing updates, you must configure the PVC or SVC to pass broadcast traffic. You can do this either per destination on the PVC or SVC by using the **broadcast** keyword at the end of an ATM mapping statement, or globally for the entire PVC or SVC using the **broadcast** command.

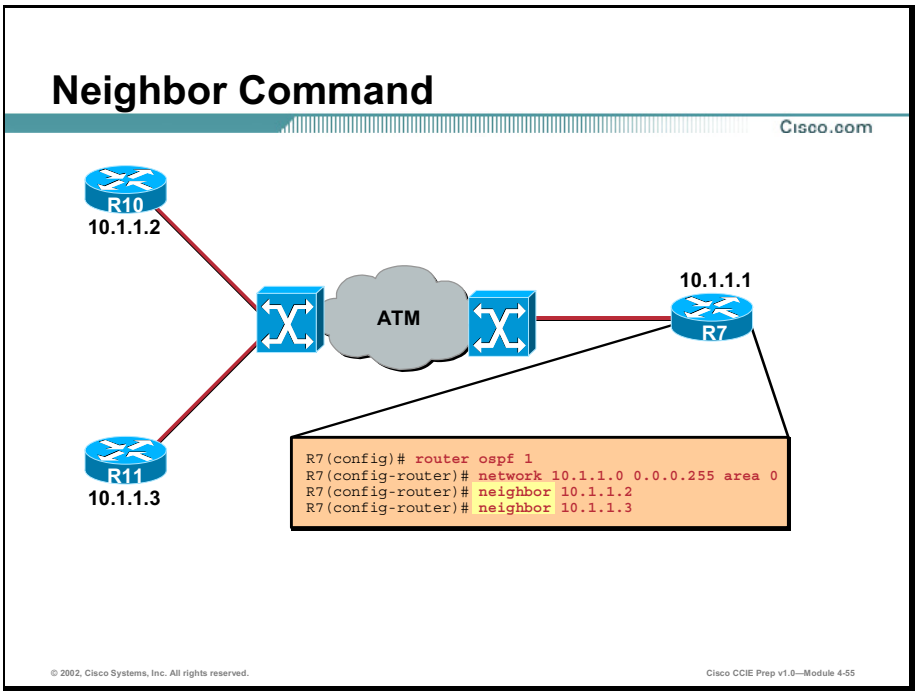

Routing protocols such as Open Shortest Path First (OSPF) and Enhanced Interior Gateway Routing Protocol (EIGRP) will require you to statically define neighbors in a non-broadcast environment. OSPF will consider ATM a non-broadcast environment by default. You can override this with the use of the **ip ospf network** command. You can statically define neighbors using the **neighbor** command within (config-router) mode for the respective routing protocol.

Configuring the AAL and Encapsulation Type

This section covers configuring the AAL and encapsulation type options.

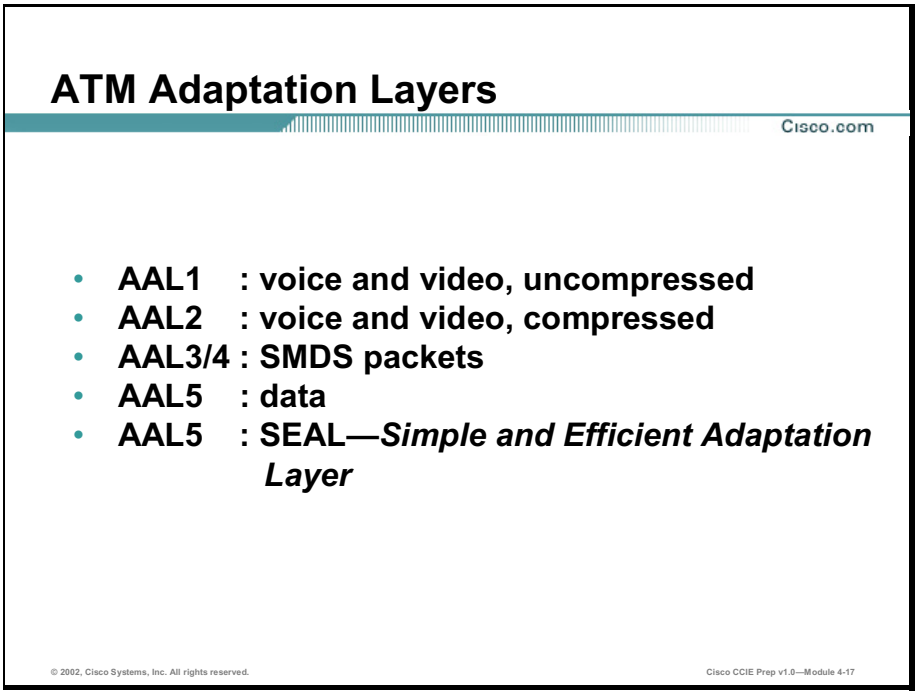

ATM Adaptation Layers: AAL1

AAL1, a connection-oriented service, is suitable for handling circuit-emulation applications, such as voice and video conferencing. Circuit-emulation service also accommodates the attachment of equipment currently using leased lines to an ATM backbone network. AAL1 requires timing synchronization between the source and destination.

ATM Adaptation Layers: AAL2

AAL2 is suitable for conveying packetized voice and video traffic.

ATM Adaptation Layers: AAL3/4

AAL3/4 supports both connection-oriented and connectionless data. AAL3/4 works primarily for network service providers and aligns closely with Switched Multimegabit Data Service (SMDS). You can use AAL3/4 to transmit SMDS packets over an ATM network.

ATM Adaptation Layers: AAL5

AAL5 is the primary AAL for data and supports both connection-oriented and connectionless data. You can use it to transfer most non-Switched Multimegabit Data Service (SMDS) data, such as classical IP over ATM and Local Area Network Emulation (LANE). Some technicians refer to AAL5 as the simple and efficient adaptation layer (SEAL), because the Segmentation and Reassembly (SAR) sublayer simply accepts the Convergence Sublayer Packet Data Unit (CS-PDU) and segments it into 48-octet SAR-PDUs without adding any additional fields.

AAL5 prepares a cell for transmission in three steps. First, the Carrier Selection (CS) sublayer appends a variable-length pad and an 8-byte trailer to a frame. The pad ensures that the resulting PDU falls on the 48-byte boundary of an ATM cell. The trailer includes the length of the frame and a 32-bit cyclic redundancy check (CRC) computed across the entire PDU. This allows the AAL5 receiving process to detect bit errors, lost cells, or cells that are out of sequence. Second, the SAR sublayer segments the CS-PDU into 48-byte blocks. The router does not add a header and trailer (as is in AAL3/4), so the router cannot interleave messages. Finally, the ATM layer places each block into the Payload field of an ATM cell. For all cells except the last, the ATM network sets a bit in the Payload Type (PT) field to zero to indicate that the cell is not the last cell in a series that represents a single frame. For the last cell, the ATM network sets the bit in the PT field to one.

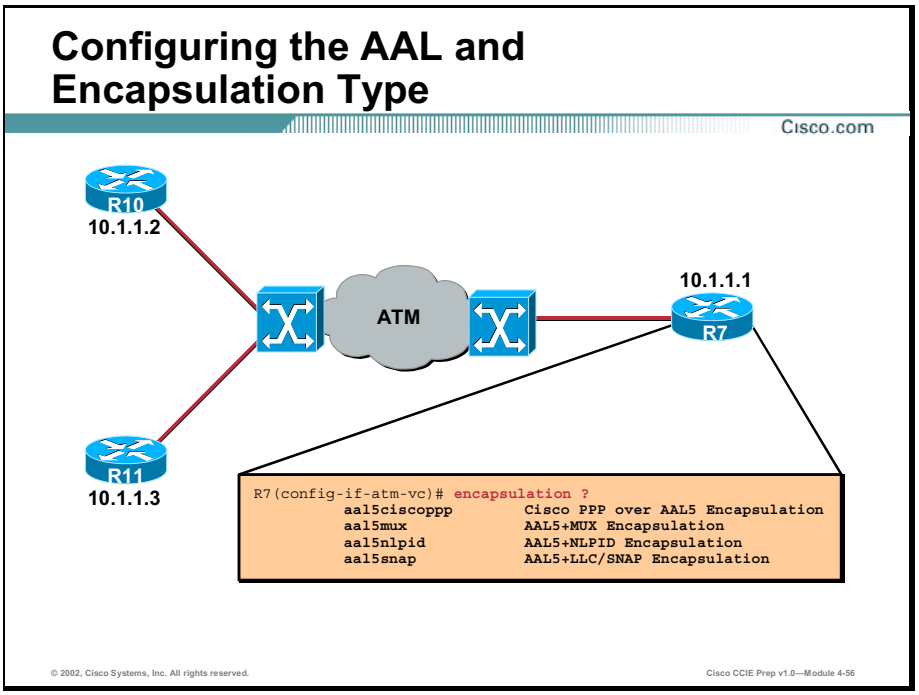

Use the **aal5mux** encapsulation option to dedicate the specified virtual circuit to a single protocol; use the **aal5snap** encapsulation option to multiplex two or more protocols over the same virtual circuit. Whether you select **aal5mux** or **aal5snap** encapsulation might depend on practical considerations, such as the type of network and the pricing offered by the network. If the network pricing depends on the number of virtual circuits, **aal5snap** might be the appropriate choice. If pricing depends on the number of bytes transmitted, **aal5mux** might be the appropriate choice, because it has slightly less overhead.

Configure the ATM adaptation layer (AAL) and encapsulation type with the command beginning in interface-ATM-VC configuration mode:

encapsulation aal5*encap*

The options for *encap* are:

```
R7(config-if-atm-vc)# encapsulation ?
```

- **aal5ciscoppp:** Cisco Point-to-Point Protocol (PPP) over AAL5 Encapsulation
- **aal5mux:** AAL5+Multiplex (MUX) Encapsulation
- **aal5nlpid:** AL5+ Network Layer Protocol ID (NLPID) Encapsulation
- **aal5snap:** AAL5+Logical Link Control/Subnetwork Access Protocol (LLC/SNAP) Encapsulation

- **ciscoppp**: For Cisco Point-to-Point Protocol (PPP) over ATM. Supported on ATM PVCs only.

- **nlpid**: Allows ATM interfaces to interoperate with High-Speed Serial Interfaces (HSSIs) that are using an ATM data service unit (ADSU) and running ATM-Data Exchange Interface (DXI). Supported on ATM PVCs only.

- **snap:** The only encapsulation supported for Inverse ARP. Logical Link Control/Subnetwork Access Protocol (LLC/SNAP) precedes the protocol datagram.

When **mux** is specified, a protocol is required:

```
R7(config-if-atm-vc)# encapsulation aal5mux ?
```

- **apollo:** Apollo Domain

- **appletalk:** AppleTalk

- **decnet:** DECnet

- **ip:** IP

- **ipx:** Novell Internetwork Packet Exchange (IPX)

- **ppp:** VC MUX PPP over AAL5 Encapsulation

- **vines:** Banyan Virtual Integrated Network Service (VINES)

- **xns:** Xerox Network Services

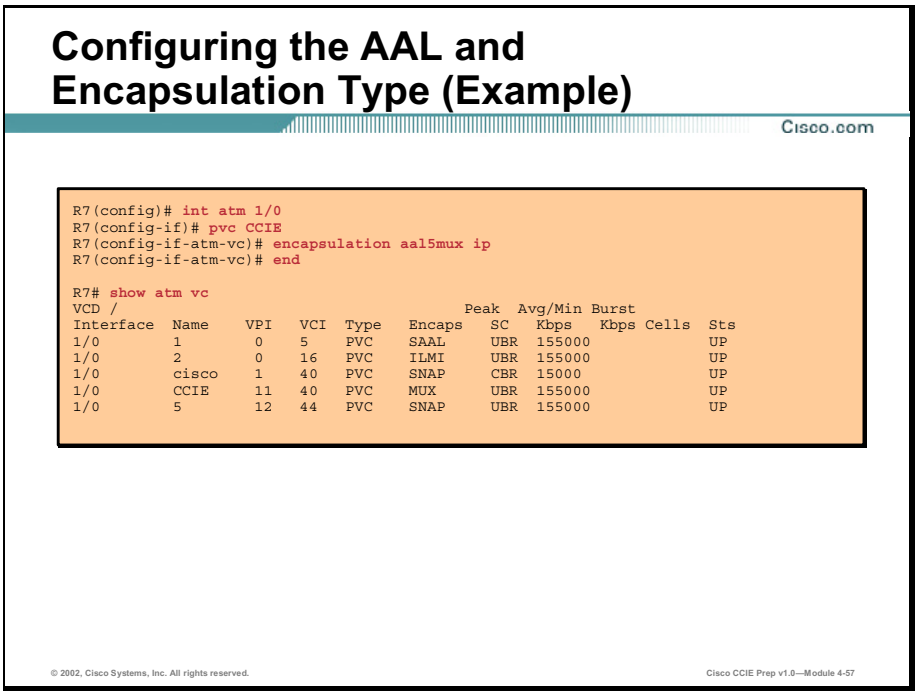

You can change the default **aal5snap** AAL and encapsulation configuration of the CCIE PVC, as seen in this example:

```
R7(config)# int atm 1/0
R7(config-if)# pvc CCIE
R7(config-if-atm-vc)# encapsulation aal5mux ip
R7(config-if-atm-vc)# end

R7# show atm vc
VCD /                                        Peak  Avg/Min Burst
Interface  Name   VPI   VCI  Type  Encaps  SC   Kbps   Kbps  Cells  Sts
1/0        1       0     5  PVC     SAAL UBR  155000              UP
1/0        2       0    16  PVC     ILMI UBR  155000              UP
1/0        cisco   1    40  PVC     SNAP CBR   15000              UP
1/0        CCIE   11    40  PVC     MUX  UBR  155000              UP
1/0        5      12    44  PVC     SNAP UBR  155000              UP
```

Configuring PVC Traffic Parameters

This section covers configuring the various PVC traffic parameters.

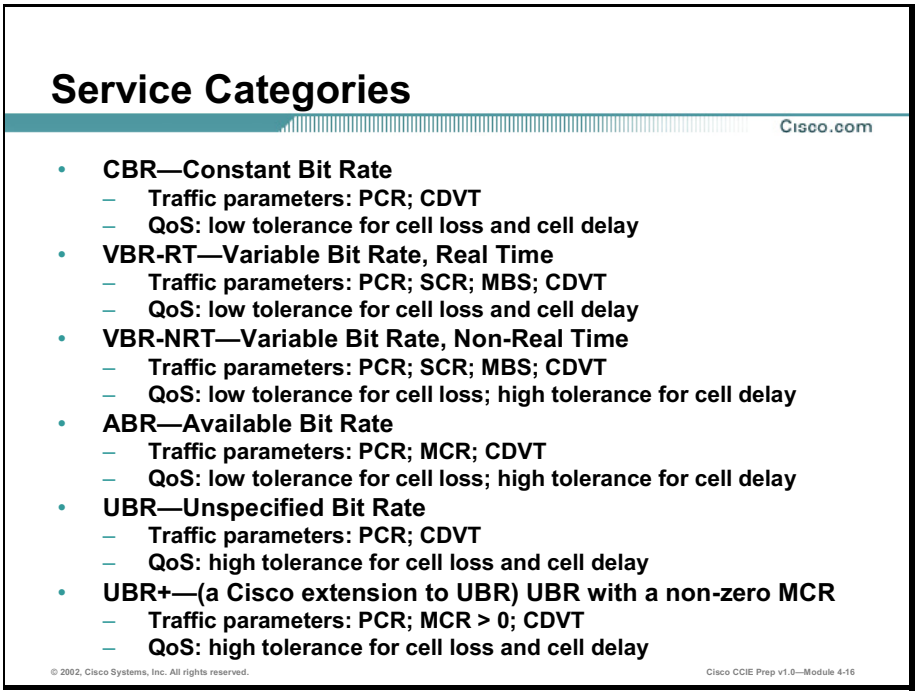

One of the main benefits of ATM is to provide distinct classes of service for the varying bandwidth, loss, and latency requirements of different applications. Some applications require constant bandwidth, while others can adapt to the available bandwidth, perhaps with some loss of quality. Still others can make use of whatever bandwidth is available and use dramatically different amounts from one instant to the next.

ATM provides five standard service categories that meet these requirements by defining individual performance characteristics, ranging from best effort (Unspecified Bit Rate [UBR]) to highly controlled, full-time bandwidth (Constant Bit Rate [CBR]). The slide shows each service category defined by the ATM Forum, along with its applicable traffic parameters and QoS characteristics.

The characteristics and uses of each service category are summarized as follows:

- **CBR** service provides constant bandwidth with a fixed timing relationship, which requires clocking synchronization. Because CBR traffic reserves a fixed amount of bandwidth, some trunk bandwidth might go unused. CBR is typically used for circuit emulation services to carry real-time voice and video.

- **Variable Bit Rate Real-Time (VBR-RT)** service provides only a partial bandwidth guarantee. Like CBR, however, some bandwidth might still go unused. Typical applications include packetized voice and video, and interactive multimedia.

- **Variable Bit Rate Non Real-Time (VBR-NRT)** service provides a partial bandwidth guarantee, but with a higher cell delay than VBR-RT. This service category is suitable for bursty applications, such as file transfers.

- **ABR** provides a best effort service, in which feedback flow control within the network is used to increase bandwidth when no congestion is present, maximizing the use of the network.

- **UBR** service provides no bandwidth guarantee, but attempts to fill bandwidth gaps with bursty data. UBR is well suited for LAN protocols, such as LAN emulation.

An additional category, **UBR+,** is a Cisco extension to UBR that provides for a nonzero Minimum Cell Rate (MCR) in the traffic contract.

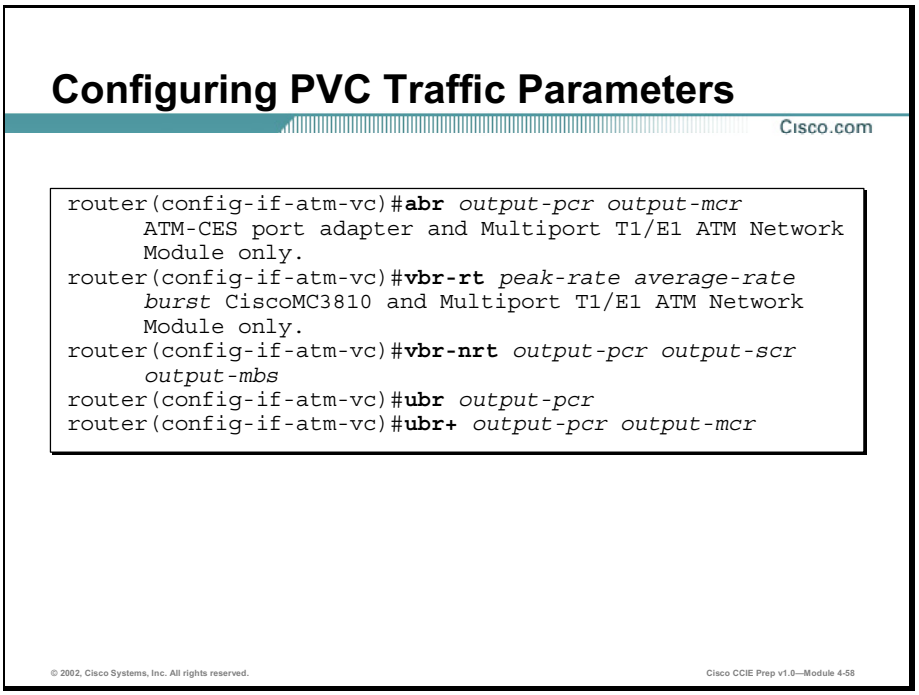

These supported traffic parameters are part of the service categories:

- Available Bit Rate (ABR)
- Variable Bit Rate Real-Time (VBR)
- Variable Bit Rate Non Real-Time (VBR-NRT)
- Unspecified Bit Rate (UBR)
- UBR+

You can specify only one of these categories per VC connection. If you enter a new service category, it replaces the existing one.

To configure VC traffic parameters, use one of the following commands in (config-if-atm-vc) mode:

- **abr:** output-pcr output-mcr

(ATM-CES port adapter and Multiport T1/E1 ATM Network Module only)

- **vbr-rt:** peak-rate average-rate burst

(CiscoMC3810 and Multiport T1/E1 ATM Network Module only)

- **vbr-nrt:** output-pcr output-scr output-mbs
- **ubr:** output-pcr

The **ubr+** *output-pcr output-mcr* command has the following parameters:

- *-pcr* = **PCR** (Peak Cell Rate)
- *-mcr* = **MCR** (Minimum Cell Rate)
- *-scr* = **SCR** (Sustained Cell Rate)
- *-mbs* = **MBS** (Maximum Burst Size)

Note Do not expect the commands in this section to work on the ATM port adapter (PA-A1 series), the ABR service class is only supported on the ATM-CES port adapter for PVCs; the 1-port ATM-25 network module only supports UBR.

For ABR VCs, you can optionally configure the amount that the cell transmission rate increases or decreases in response to flow control information from the network or destination. You will use the following command, in (config-if-atm-vc) mode:

atm abr rate-factor [*rate-increase-factor*] [*rate-decrease-factor*]

You will feed in the ABR rate factor: the default increase and decrease rate factors is 1/16.

This is an example for an ABR PVC:

```
R7(config)# interface atm 1/0
R7(config-if)# pvc CCIE
R7(config-if-atm-vc)# abr ?
<64-155000>  Peak Cell Rate(PCR) in Kbps

R7(config-if-atm-vc)# abr 640 ?
<0-640>  Minimum Cell Rate(MCR) in Kbps

R7(config-if-atm-vc)# abr 640 64 ?
<cr>

R7(config-if-atm-vc)# abr 640 64
R7(config-if-atm-vc)# end
```

```
R7# show atm pvc CCIE
ATM1/0: VCD: 4, VPI: 11, VCI: 40, Connection Name: CCIE
ABR, PeakRate: 640, Minimum Rate: 64, Initial Rate: 64, Current   Rate: 0

R7# show atm vc
VCD /                                        Peak  Avg/Min Burst
Interface  Name   VPI  VCI Type  Encaps  SC  Kbps  Kbps  Cells  Sts
1/0        1       0    5  PVC    SAAL UBR  155000              UP
1/0        2       0   16  PVC    ILMI UBR  155000              UP
1/0        cisco   1   40  PVC    SNAP CBR   15000              UP
1/0        CCIE   11   40  PVC    MUX  ABR     640  64          UP
1/0        5      12   44  PVC    SNAP UBR  155000              UP
```

Troubleshooting ATM

This section covers the various show and debug commands that are available to troubleshoot and verify ATM configurations.

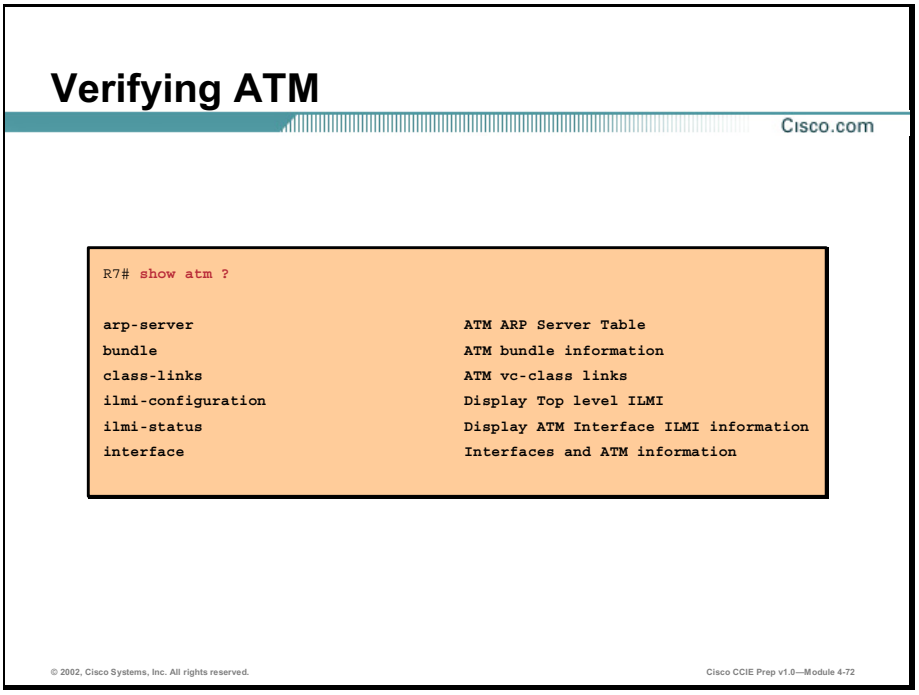

The **show atm** command syntax is useful during troubleshooting and shows you if there is a problem in the server configuration.

```
R7# show atm ?
```

- **arp-server:** ATM ARP Server Table
- **bundle:** ATM bundle information
- **class-links:** ATM vc-class links
- **ilmi-configuration:** Display Top level ILMI
- **ilmi-status:** Display ATM Interface ILMI information

This is very useful in case you misconfigured the ILMI PVC. It tells you there is no communication with the ATM switch, which in turns means no ATM NSAP on-the-fly configuration and no SVCs.

- **interface:** Interfaces and ATM information

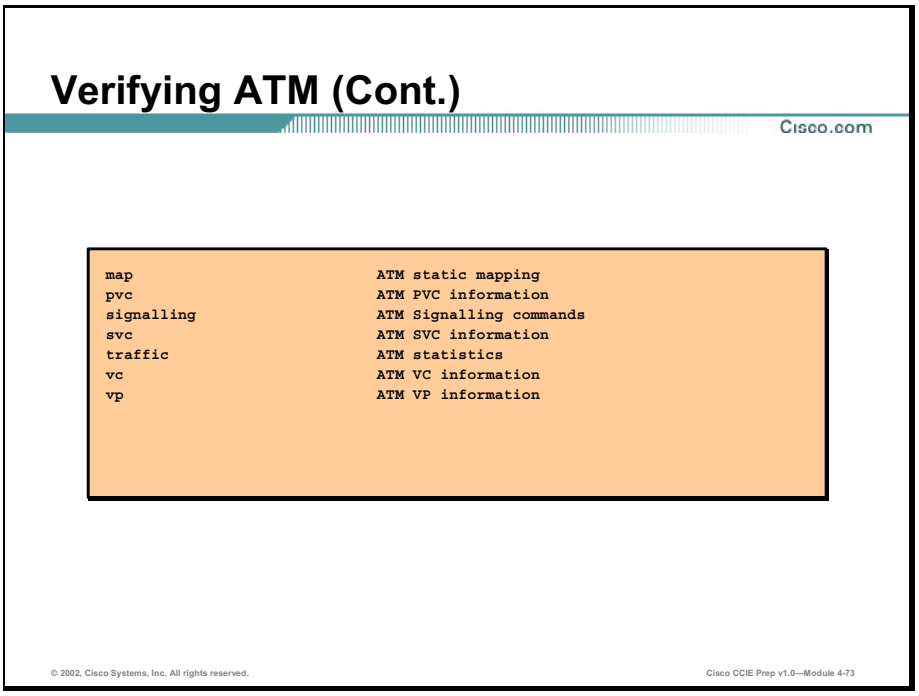

The **show atm map** command is very useful for checking any kind of virtual connection that transports higher-level protocols. Whenever you lose or cannot achieve connectivity, this is one of the first commands to use to check the configuration.

- **map:** ATM static mapping

show atm pvc, as well as the related commands, **show atm svc**, **show atm vc**, **show atm vp**, give information on the connections you have configured. It is a good idea to use them every time you create a new connection, just to double check on the data you have introduced into the router.

- **pvc:** ATM PVC information

show atm signalling statistics informs you about the signaling process.

- **signalling:** ATM Signaling commands

Note	Cisco's IOS uses a spelling of "**signalling**" for the word "signaling."

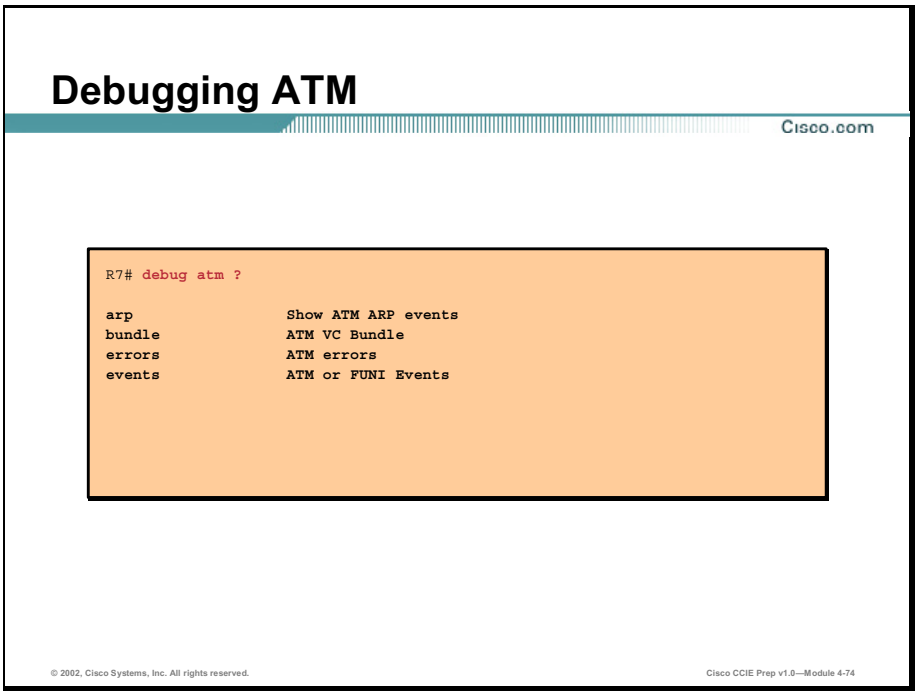

- **arp:** Show ATM ARP events. Displays messages pertaining to the process of mapping an ATM Network Service Access Point (NSAP) address to an IP address.

Here is a sample output for this command:

```
R7# debug atm arp
ATM ARP events debugging is on
R7# ping 10.1.
20:32:54: ATMARP(ATM1/0.2): Learned address through INARP reply for CCIE
R7# ping 10.1.1
20:32:57: ATMARP(ATM1/0.2): Learned address through INARP reply for CCIE
R7# ping 10.1.1.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 4/4/4 ms
```

The **error** command is shown here.

- **errors:** ATM errors – Displays error that occur for the ATM process on this router. This can include errors on an ATM interface or an error that involves any ATM related activity.

Here is a sample output for this command:

```
R7# debug atm errors
ATM errors debugging is on
R7#
20:35:21: ATM(ATM1/0.7) Send:Error in encapsulation, No VC for address
0xE000000A
20:35:22: ATM(ATM1/0.2) Send:Error in encapsulation, No VC for address
0xE000000A
```

The **event** command is shown here.

- **events:** ATM or Frame-based User to Network Interface (FUNI) Events. This command displays ATM events that occur on the ATM interface processor and is useful for diagnosing problems in an ATM network. It provides an overall picture of the stability of the network. In a stable network, the **debug atm events** command does not return any information. If the command generates numerous messages, the messages can indicate the possible source of problems.

Here is a sample output for this command:

```
R7# debug atm events
ATM events debugging is on
R7# ping 10.1.1.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 4/4/4 ms
20:39:52:  Reserved bw for 0/0 Available bw = 155000
20:39:52: rs8234_setup_vc(ATM1/0): vc:900 vpi:0 vci:74
20:39:52: rs8234_setup_vc_common() VCD=900 vp/vc=0/74 etype=0
20:39:52: rs8234_setup_cos(ATM1/0): vc:900 wred_name:- max_q:0
20:39:52:  Created 64-bit VC counters
```

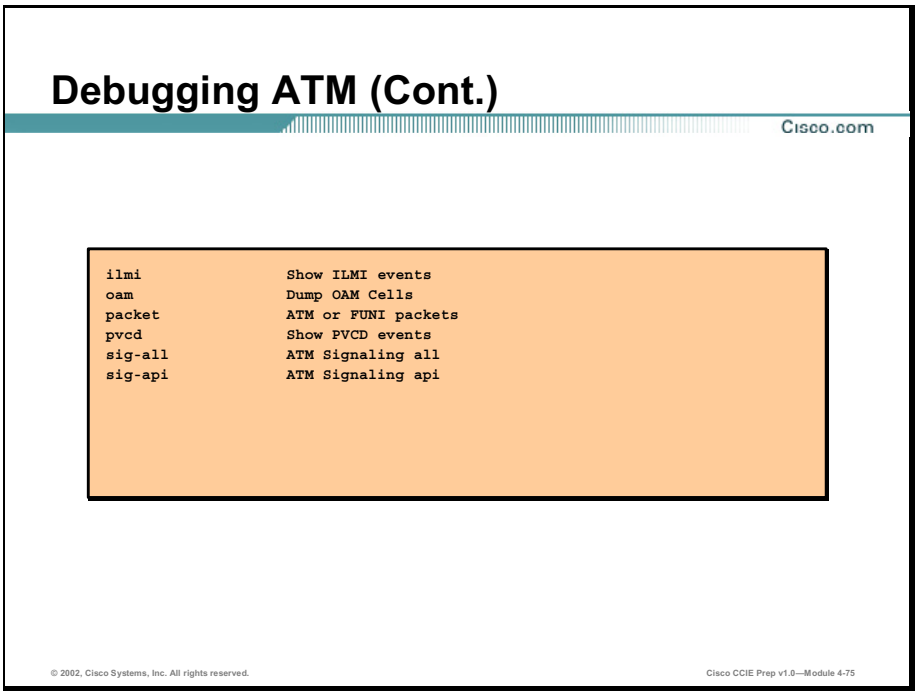

The following command will assist you in showing ILMI events.

- **ilmi:** Show ILMI events

Sample output (after a **shutdown**-**no shutdown** on the main ATM interface) for this command is shown here:

```
20:44:52: ILMI(ATM1/0):Response received for request 2042
20:44:52: ILMI Continuing Getnext processing
20:44:52: ILMI: Delivering GetNext response to Client: atmfVccEntry.13.0.1.102-1
20:44:52: ILMI(ATM1/0):Sending out Request 2043
20:44:52: ILMI(ATM1/0):Response received for request 2043
20:44:52: ILMI Continuing Getnext processing
20:44:52: ILMI: Delivering GetNext response to Client: atmfVccEntry.13.0.1.103-1
20:45:33: %SYS-3-MSGLOST: 5 messages lost because of queue overflow
20:44:52: ILMI(ATM1/0):Sending out Request 2044
20:44:52: ILMI(ATM1/0):Response received for request 2044
20:44:52: ILMI Continuing Getnext processing
20:44:52: ILMI: Delivering GetNext response to Client: atmfVccEntry.13.0.9.912-1
```

Here are some other keywords that are available for the **debug atm** command.

- **packet:** ATM or FUNI packets. The **debug atm packet** command displays all process-level ATM packets for both outbound and inbound packets. This command is useful for determining whether packets are being received and transmitted correctly.

- **pvcd:** Show PVCD events. Displays the PVC Discovery events and ILMI MIB traffic used when discovering PVCs.

The signaling related **debug atm** commands could help you to solve SVC-related issues:

- **sig-all:** ATM Signaling all. Displays all ATM signaling activity. This includes events and errors.

Summary

This section summarizes the key points discussed in this lesson.

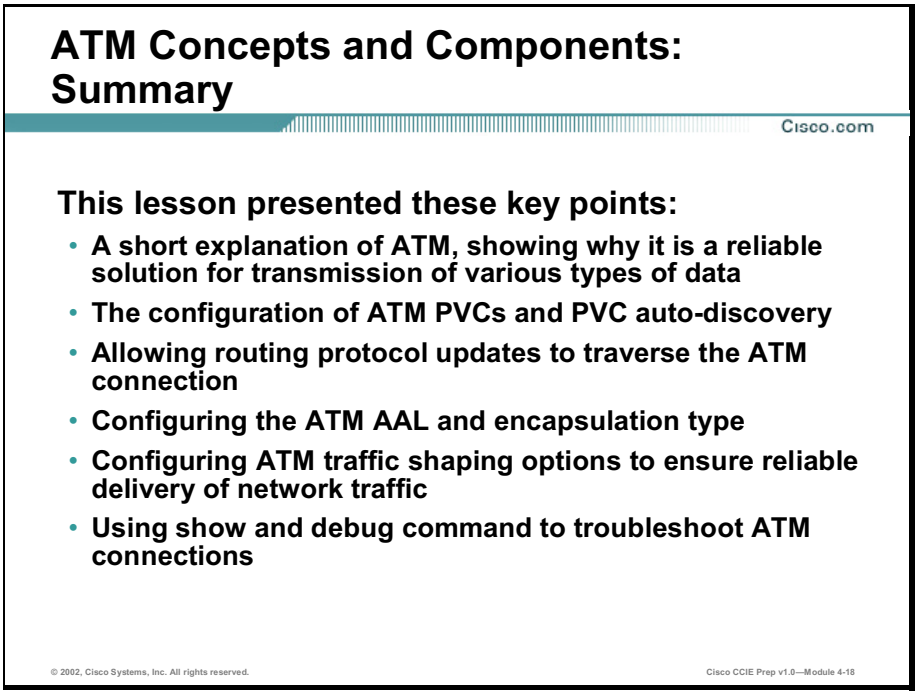

Next Steps

After completing this lesson, go to:

- Verifying ATM Configuration

References

For additional information, refer to these resources:

- http://www.atmforum.com/

Lesson Assessment

1. ATM networks is closely related to which network type?

   Q1) Synchronous

   Q2) Asynchronous

   Q3) Dedicated

   Q4) None of the above

2. Which of the following steps are REQUIRED to configure an ATM connection? (choose two)

   Q1) Create a PVC

   Q2) Map a protocol address to a PVC

   Q3) Configure the AAL and encapsulation type

   Q4) Configure PVC traffic parameters

3. Configuring ILMI on an ATM connection allows it to discover which type of address?

   Q1) Network layer

   Q2) VPI/VCI

   Q3) DLCI

   Q4) Session layer

4. Which AAL encapsulation type would you use if you would like to run multiple protocols over a single ATM VC?

   Q1) Aal5snap

   Q2) Aal5mux

   Q3) Aa5encap

   Q4) None of the above

3

ISDN Technologies

Overview

Integrated Services Digital Network (ISDN) is still used in business markets because it allows multiple digital channels to operate simultaneously over a single circuit. It can support voice, data, and video over existing phone wiring. This module examines the configuration of Dial-On-Demand Routing (DDR) over an ISDN Basic Rate Interface (BRI) link, the configuration and features of Point-to-Point Protocol (PPP), using ISDN DDR as a backup link, and the Cisco Internetwork Operating System (IOS) tools available to verify correct ISDN network operation.

Upon completing this module, you will be able to:

- Configure ISDN using physical interfaces
- Configure ISDN using Dialer Profiles
- Configure PPP and utilize its advanced features
- Use ISDN as a backup connection
- Troubleshoot ISDN connectivity

Outline

The module contains these lessons:

- ISDN Configuration
- PPP Features
- Using ISDN as a Backup Connection
- Troubleshooting

ISDN Configuration

Overview

This lesson reviews the basic functionality of Integrated Services Digital Network (ISDN) as a Dial-on-Demand (DDR) Wide Area Network (WAN) connection. This lesson also covers the configuration of differences between Legacy DDR and Dialer Profiles.

Importance

ISDN is a key technology in the Cisco Certified Internetwork Expert (CCIE) lab. Knowing how to configure an ISDN interface to work as a Dial-on-Demand Routing (DDR) circuit will be the basis for many subsequent tasks.

Objectives

Upon completing this lesson, you will be able to:

- Describe ISDN Functionality
- Configure Dial-on-Demand Routing (DDR)
- Configure Dialer Profiles

Learner Skills and Knowledge

To fully benefit from this lesson, you must have these prerequisite skills and knowledge:

- Passed the CCIE written qualification exam
- Completed the Building Cisco Remote Access Networks (BCRAN) course or have the equivalent knowledge

Outline

This lesson includes these sections:

- Overview
- Network Diagram
- Basic Configuration
- Dial-on-Demand Routing (DDR)
- Dialer Profiles
- Summary
- Lesson Assessment (Quiz)

Network Diagram

This network diagram will be the basis for ISDN configuration in this course.

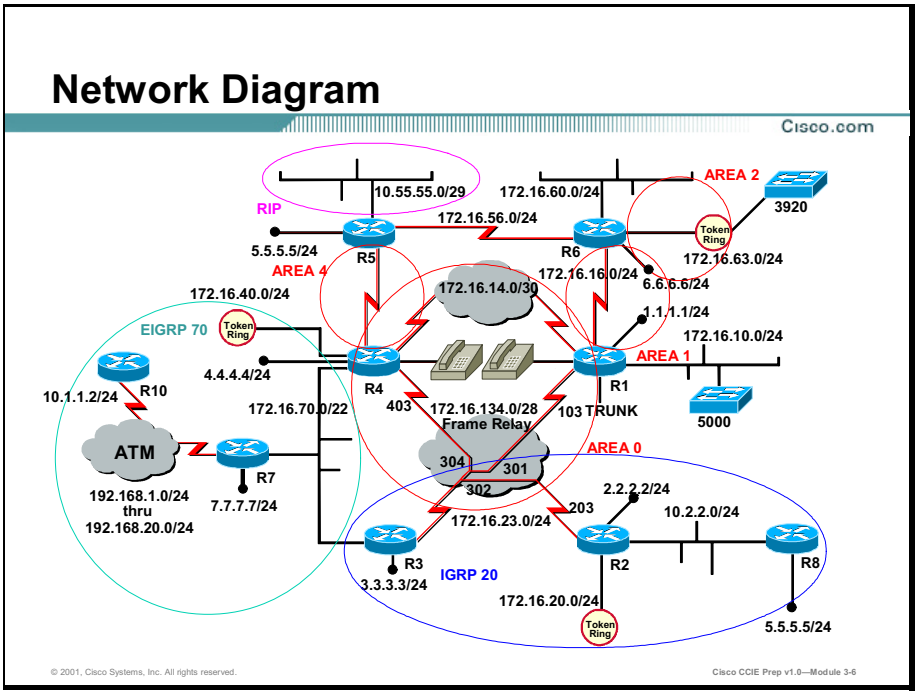

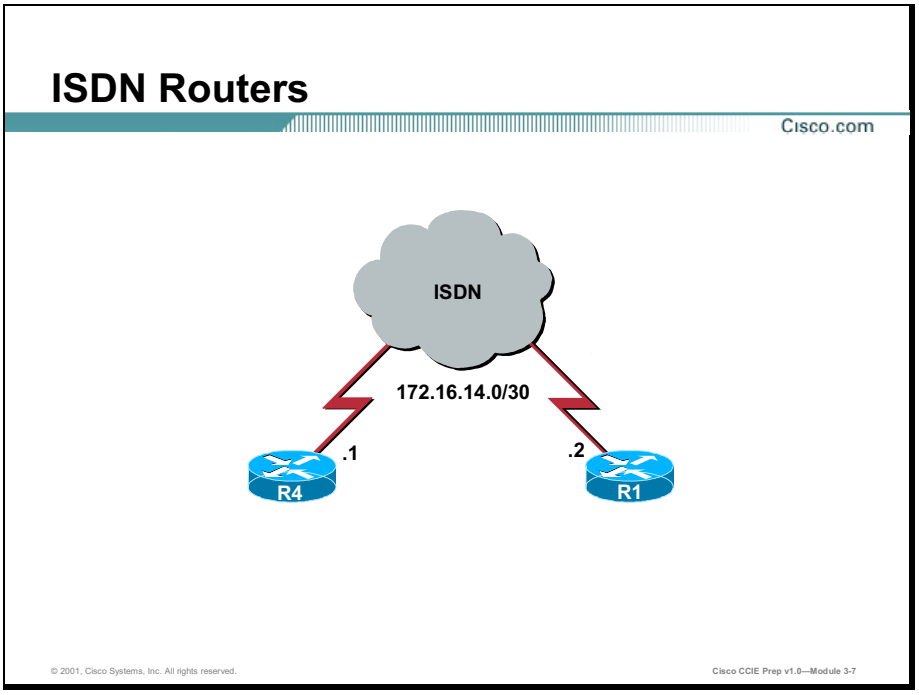

Basic Configuration

Basic configuration of ISDN involves setting up DDR on the physical interfaces on both sides of the ISDN connection.

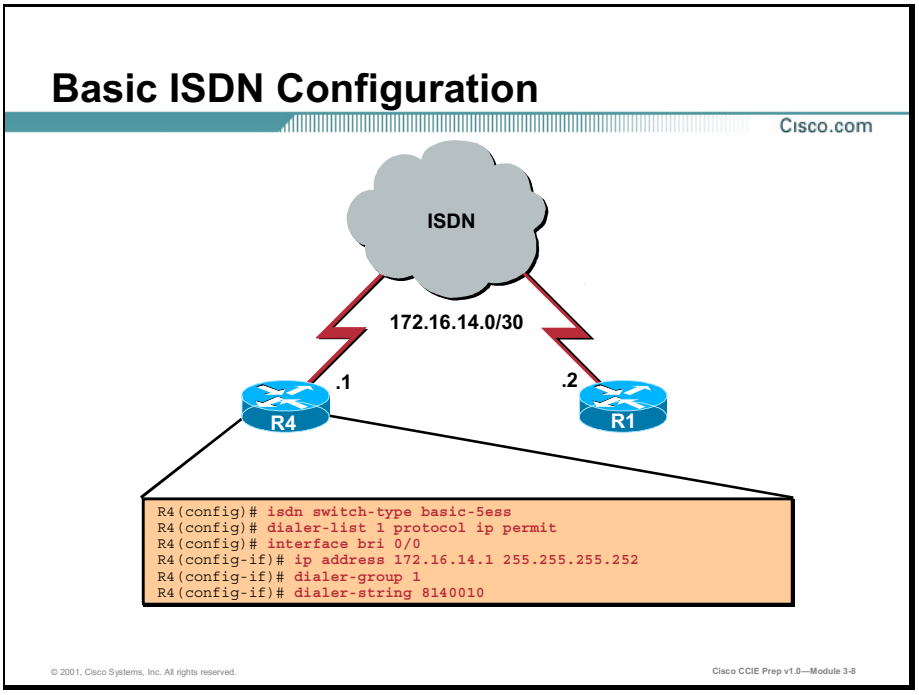

Shown here is the most basic ISDN configuration; it has limited capabilities, uses bandwidth poorly, and does not scale well. High-Level Data Link Control (HDLC) will be used as the encapsulation because an encapsulation type has not been specified. Therefore, both B channels will be used in each direction: one channel for sending data and the other for receiving data, which means that dialer strings are required on both sides of the link.

Dial-on-Demand Routing (DDR)

Dial-on-Demand Routing (DDR) allows an ISDN interface to be brought up only when certain traffic needs to cross the link. This allows ISDN connectivity to be established on an as-needed basis, reducing the costs associated with ISDN connectivity.

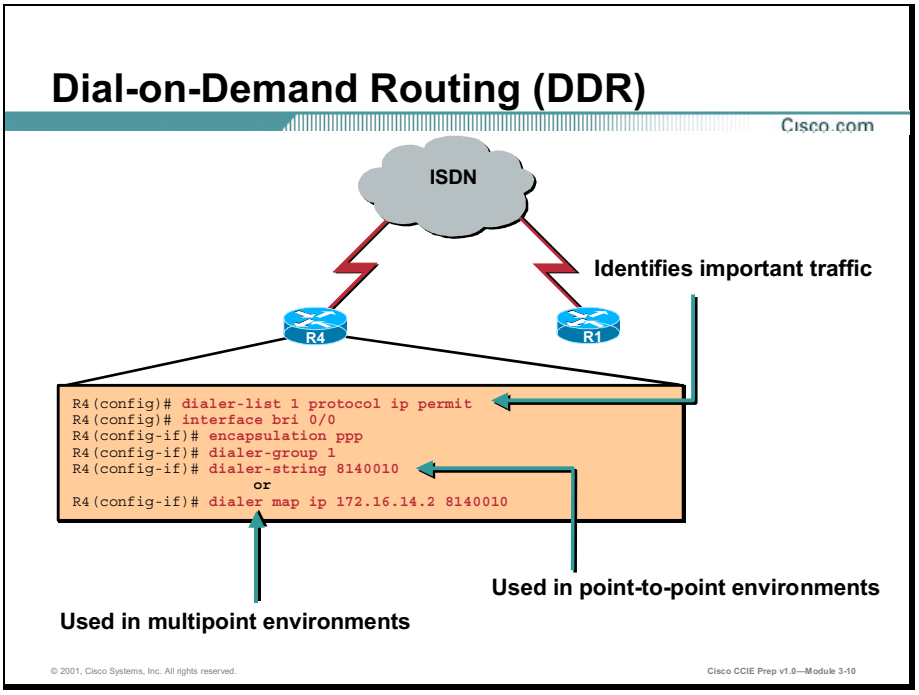

The DDR configuration shown here will allow you to take advantage of the many properties of Point-to-Point Protocol (PPP). The interesting traffic has been identified as any IP type traffic. The **dialer string** command should only be used in a point-to-point environment. If you are configuring ISDN in a point-to-multipoint environment, you will most often be using the **dialer map** command.

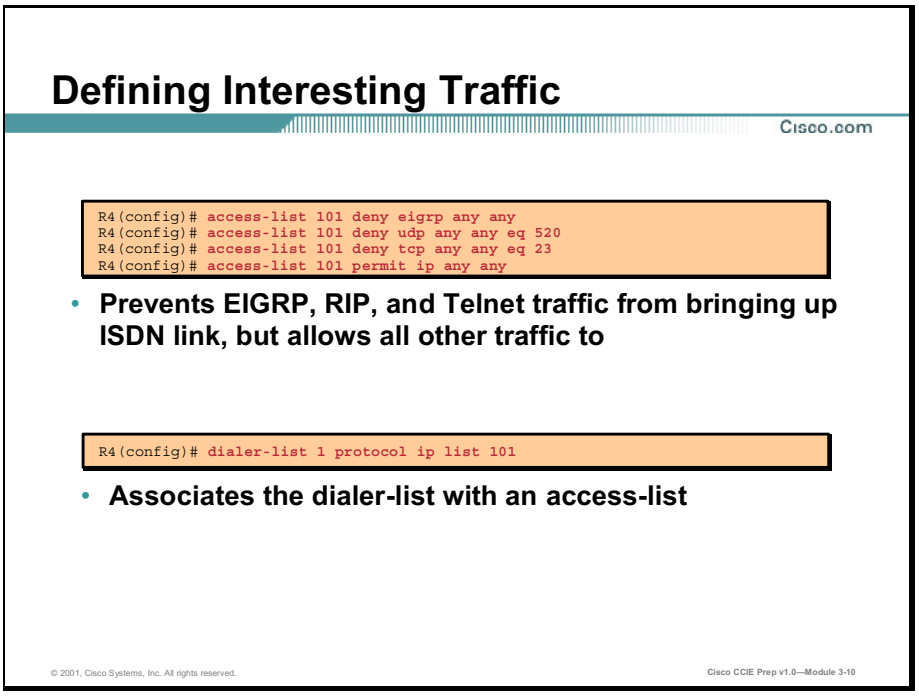

The **dialer-list protocol** form of the **dialer-list** command defines interesting traffic based on protocol. The **dialer-list protocol** <*protocol*> **list** form of this command allows for a more granular definition of interesting traffic using an access list.

In the example shown, Enhanced Interior Gateway Routing Protocol (EIGRP) routing protocol updates, Routing Information Protocol (RIP) routing updates, and Telnet traffic are not classified as interesting traffic and therefore will not initiate calls on the ISDN circuit.

To complete the DDR configuration, apply the dialer-list to an ISDN interface with the **dialer-group** command.

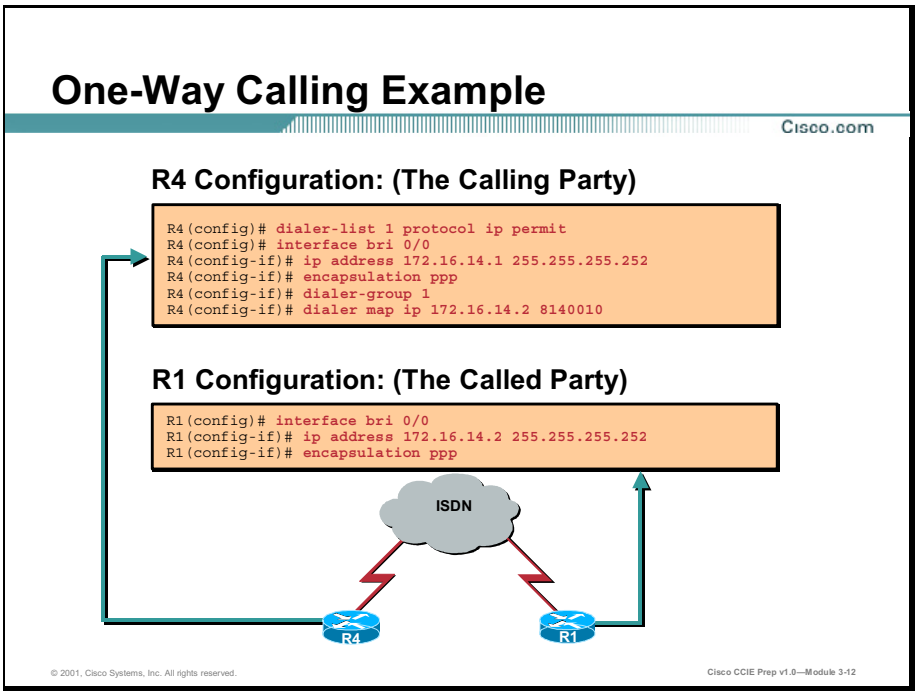

Suppose you only want R4 to initiate calls. In this case, you can simply remove any dialer strings or dialer maps from R1's configuration. When R4 initiates a call to R1, a dynamic ISDN mapping will occur for return traffic.

Also, since you will never initiate a call from R1, you can also remove the interesting traffic parameters (dialer-list and dialer-group) from R1's configuration.

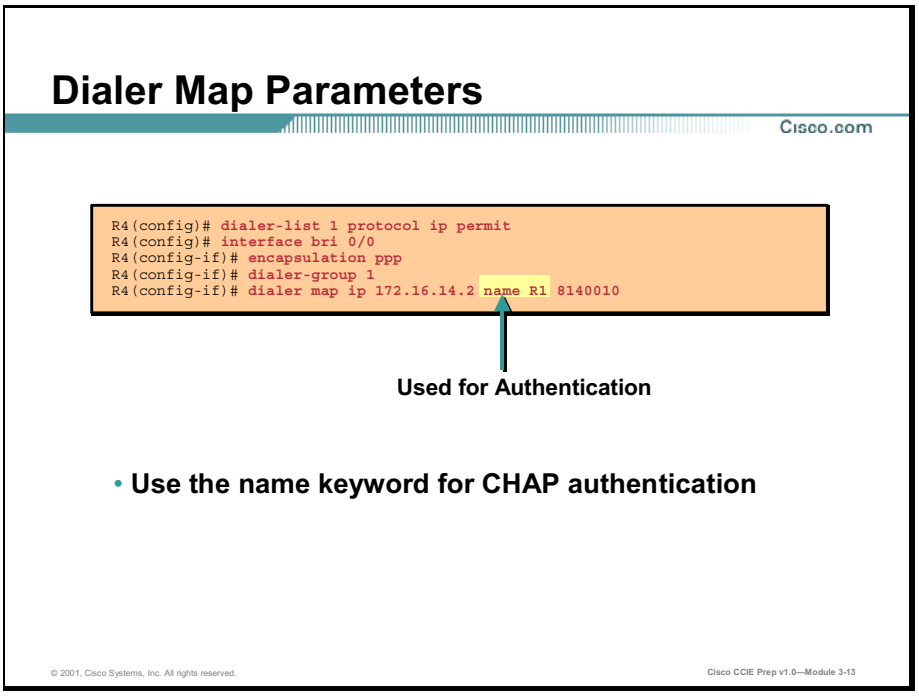

A large benefit of using PPP is its ability to perform secure authentication via Challenge-Handshake Authentication Protocol (CHAP). In order to perform PPP authentication, the **name** keyword should be included in the dialer map. In a point-to-point ISDN environment using dialer strings instead of dialer maps, the equivalent to the **name** keyword is the **dialer remote-name** command.

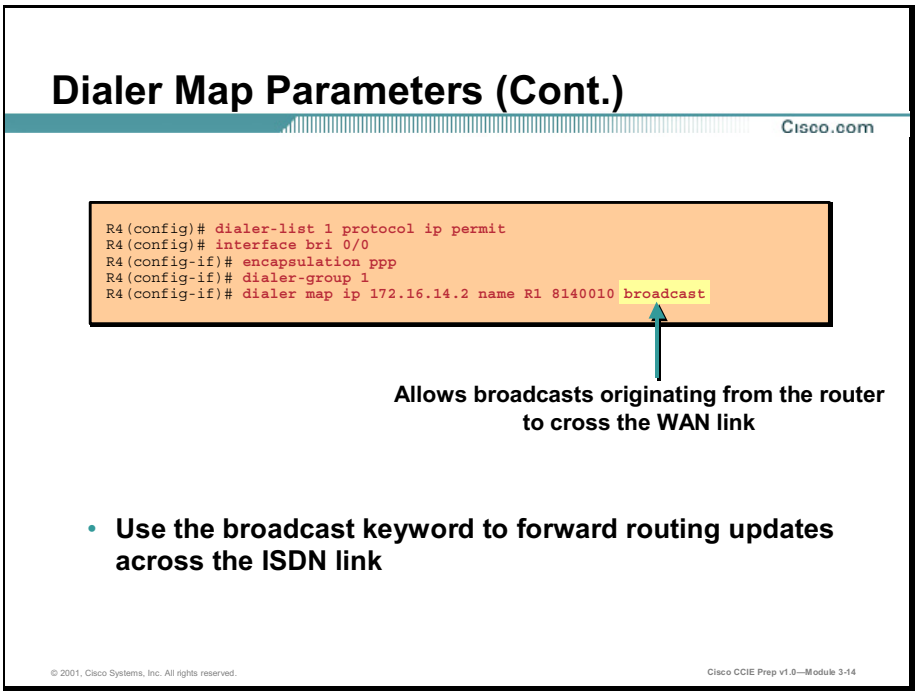

In this example, static routing over the ISDN link has been removed in favor of using a dynamic routing protocol like Routing Information Protocol (RIP) or Open Shortest Path First (OSPF). RIP sends routing updates to UDP port 520 every 30 seconds by default and OSPF sends multicast packets to 224.0.0.5. The basic dialer map statement is insufficient because broadcast and multicast traffic will not be sent across the link. To allow these traffic types to be sent over the ISDN link, use the **broadcast** keyword in the dialer map statement.

The basic dialer map statement is insufficient because broadcast and multicast traffic will not be sent across the link. To allow these traffic types to be sent over the ISDN link, use the broadcast keyword in the dialer map statement.

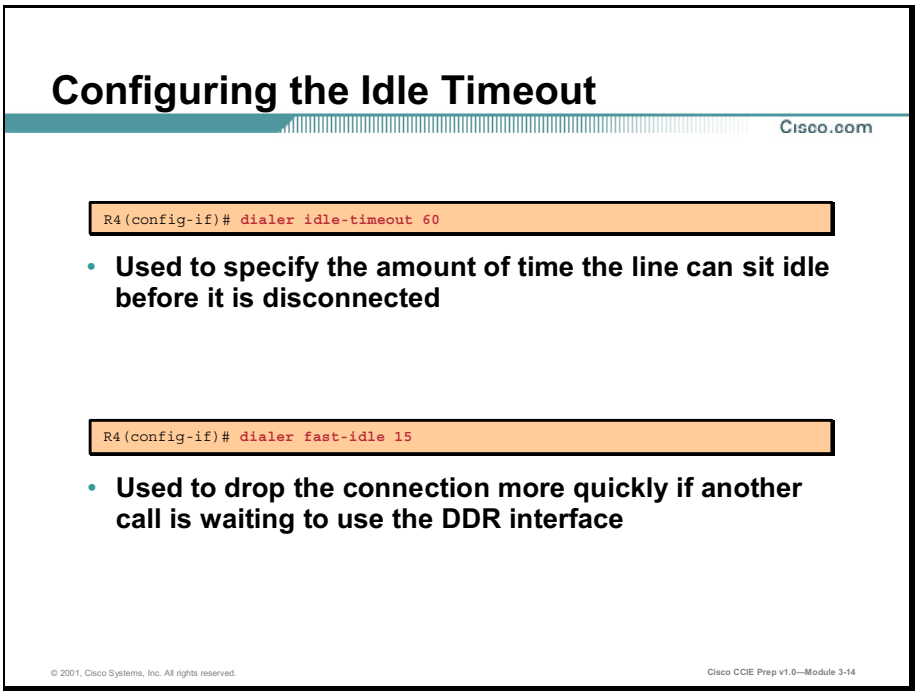

To specify the idle time before the line is disconnected, use the **dialer idle-timeout** command in interface configuration mode. Remember the dialer idle-timeout command is based on interesting traffic defined in the dialer-list command. Use the **no** form of this command to reset the idle timeout to the default value of 120 seconds.

If both your ISDN Basic Rate Interface (BRI) channels are being used for calls to two different locations, and a call to a different branch office has been requested, the router will have to wait until the idle timeout has been reached on one of the first two channels before it can place the new call. You can use a different timer when contention for a B channel exists. That timer is referred to as the fast idle timer. If this timer has been defined, when contention for a B channel exists, instead of waiting for the idle-timeout to reach zero, this shorter timeout is used to drop the line faster so that it can be used for the newly queued calls. You configure the fast idle time with the **dialer fast-idle** command in interface configuration mode. Use the **no** form of this command to return it to the default value of 20 seconds.

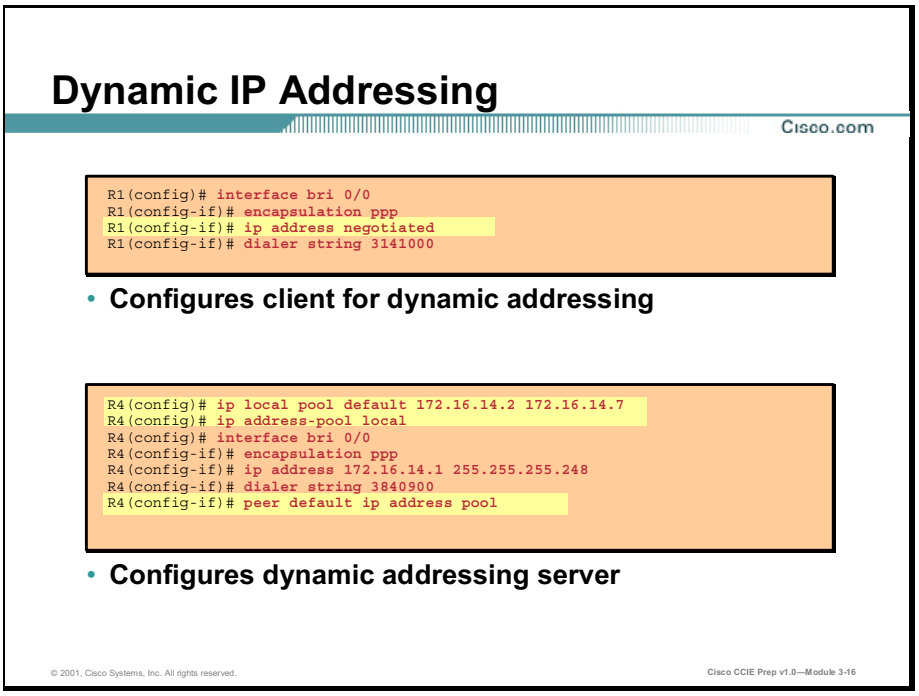

When using PPP encapsulation, you can have the (Internet Protocol) IP address of the client negotiate its IP address from the server (call initiator). Using dynamic IP address negotiation (PPP/IPCP) at each remote site substantially reduces Internet access costs. Static IP addresses are considerably more expensive to purchase than are dynamically allocated IP addresses. IP address negotiation is usually performed in a hub and spoke fashion where the initiator can be either the hub or the server. This allows you to apply policies to these known IP addresses on the server side, as well as have a streamlined, consistent configuration on your spoke routers.

In this example, R4 has created a default pool of IP addresses in the range 172.16.14.2-172.16.14.7. Next, specify that the local pool of IP addresses will be obtained from this default pool. Finally, specify that the peer will obtain an IP address from this default pool of IP addresses. In this scenario, R4 must initiate the call to R1, as R1 will not have an IP address until the IPCP negotiations have completed. This would normally be accomplished through the use of static routes.

Dialer Profiles

Dialer Profiles allow a great deal of flexibility in the configuration of DDR circuits, allowing one physical interface to be used for multiple functions.

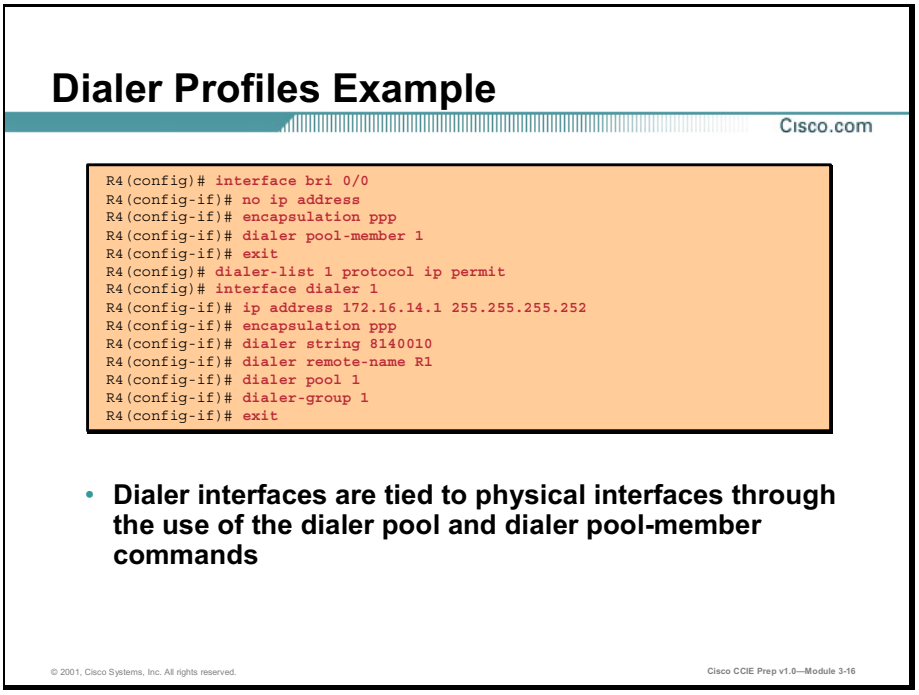

Shown here is a configuration that will be examined over the following pages. In this example dialer interface 1 will be used to backup serial 0/0 and dialer interface 2 will be used to transfer e-mail to the Central Office.

```
R4(config)# interface bri 0/0
R4(config-if)# no ip address
R4(config-if)# encapsulation ppp
R4(config-if)# dialer pool-member 1
R4(config-if)# exit
R4(config)# dialer-list 1 protocol ip permit
R4(config)# interface dialer 1
R4(config-if)# ip address 172.16.14.1 255.255.255.252
R4(config-if)# encapsulation ppp
R4(config-if)# dialer string 8140010
R4(config-if)# dialer remote-name R1
R4(config-if)# dialer pool 1
R4(config-if)# dialer-group 1
R4(config-if)# exit
R4(config)# dialer-list 2 protocol ip list 102
R4(config)# access-list 102 permit tcp any any eq smtp
```

```
R4(config)# interface dialer 2
R4(config-if)# ip address 172.16.30.1 255.255.255.252
R4(config-if)# encapsulation ppp
R4(config-if)# dialer string 9650001
R4(config-if)# dialer remote-name CentralOffice
R4(config-if)# dialer pool 1
R4(config-if)# dialer-group 2
R4(config-if)# exit
R4(config)# interface serial 0/0
R4(config-if)# backup interface dialer 1
```

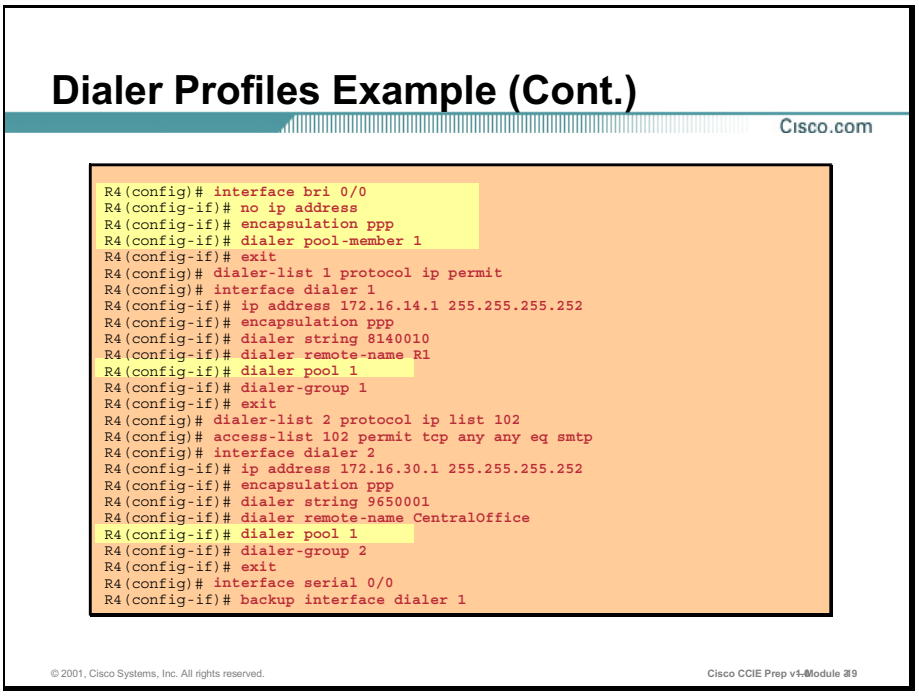

First, all Legacy DDR commands are removed from the physical interface. No IP address has been assigned to the physical interface, the IP address will be obtained when the profile is mapped to the physical interface. Mapping is performed with the commands **dialer pool-member 1** and **dialer pool 1**. The single physical interface stipulates that any logical interface that wishes to use this interface must be a member of dialer pool 1. Both dialer interfaces are configured to use BRI 0/0, as they are members of dialer pool 1.

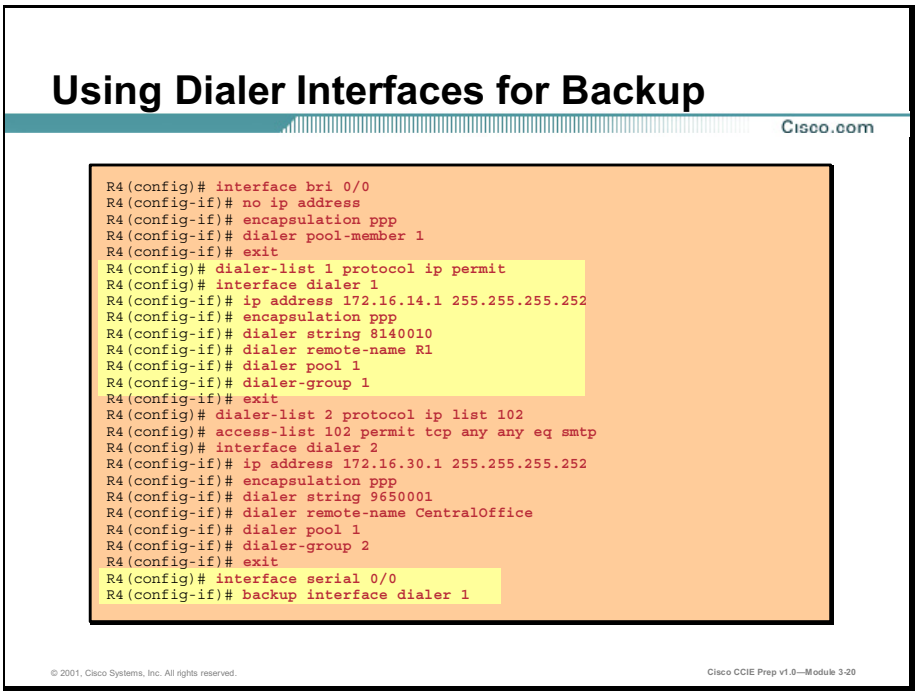

Next, a logical dialer interface is created. This interface will be used as a backup interface to serial 0/0. An IP address and a dialer string are assigned to the dialer interface. Interesting traffic is defined that will trigger this dialer interface. In this case, any IP traffic will bring up the link. This logical interface was specified as a member of dialer pool 1 to map it to the physical BRI 0/0 interface.

Notice that a new command has been issued on this dialer interface **dialer remote-name**. This command specifies the name of the device that this interface wishes to call. In this case, the opposite device has a name of R1. This remote-name is used during CHAP authentication.

In this example, you only want the dialer 1 interface to activate when your Frame Relay connection drops. To accomplish this task, dialer 1 is defined as a backup interface for serial 0/0. This configuration will allow you to have a redundant backup connection, but still use the ISDN circuit for DDR when the primary interface is functioning correctly. This configuration is impossible with Legacy DDR.

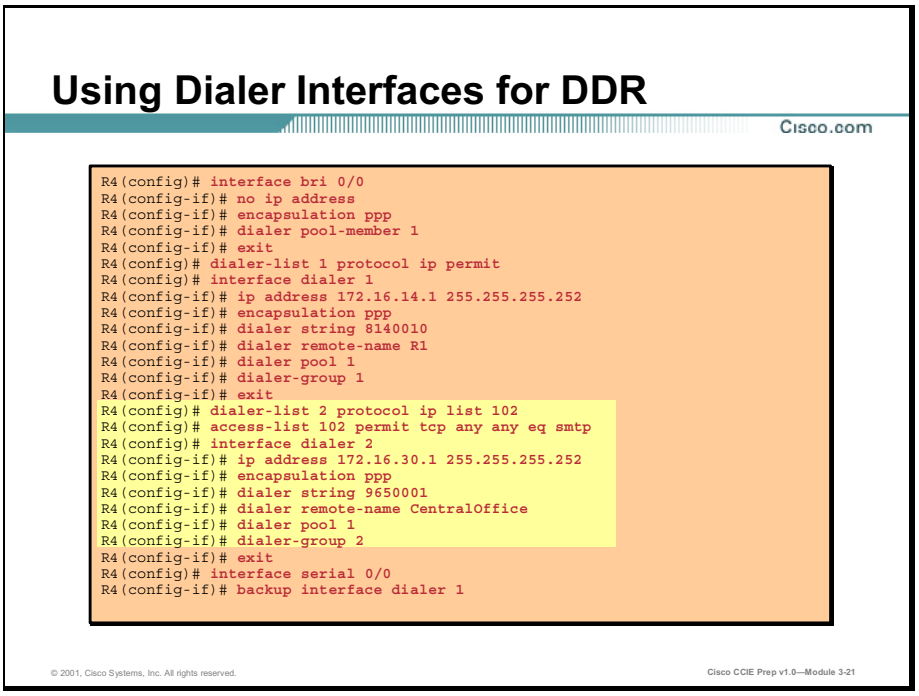

The second goal was to allow e-mail transfer to the Central Office. This goal is accomplished with the use of a second dialer interface. Here the only traffic allowed to bring up the link (interesting traffic) is Simple Mail Transfer Protocol (SMTP), which is used to send e-mail. Assign an IP address to this interface, as well as the dialer string used to reach the destination, which has a remote name of CentralOffice. Finally, specify this dialer interface to be a member of pool 1 to use the physical BRI 0/0 interface when needed.

Dialer profiles are extremely versatile and can help accomplish goals not ordinarily obtained through the use of Legacy DDR. Another advantage is that dialer profiles are easy to configure and implement.

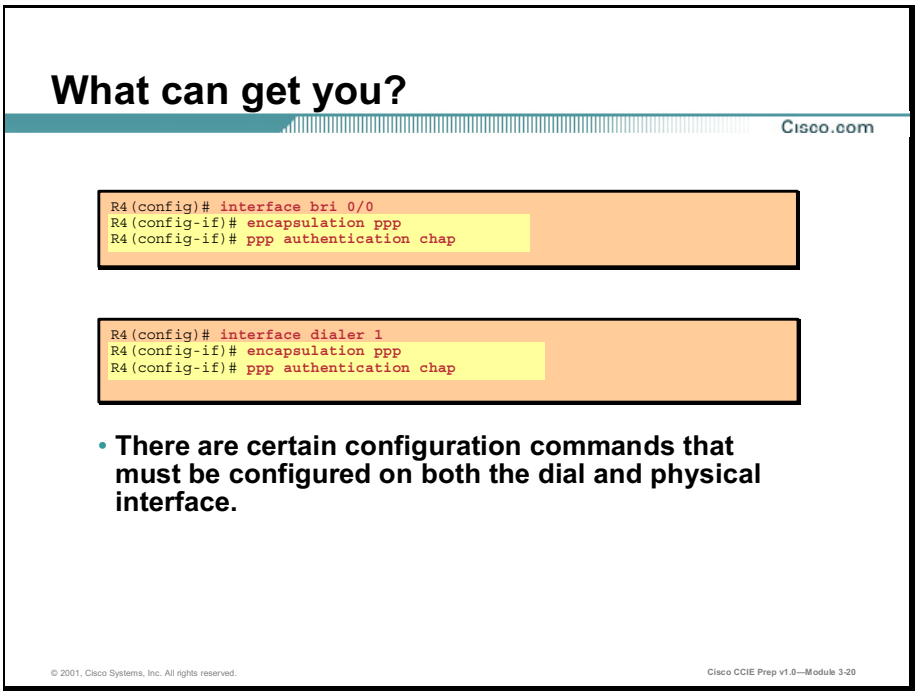

When using dialer profiles there are certain configuration commands that must be configured on both the physical interface, as well as the logical interface. In the previous examples, the command **encapsulation ppp** was used on the physical, as well as the dialer interfaces. If you perform PPP authentication, such as CHAP authentication, the command **ppp authentication chap** is required on both the physical and the logical interfaces.

```
R4(config)# interface bri 0/0
R4(config-if)# no ip address
R4(config-if)# encapsulation ppp
R4(config-if)# ppp authentication chap
R4(config-if)# dialer pool-member 1
R4(config-if)# exit
R4(config)# dialer-list 1 protocol ip permit
R4(config)# interface dialer 1
R4(config-if)# ip address 172.16.14.1 255.255.255.252
R4(config-if)# encapsulation ppp
R4(config-if)# ppp authentication chap
R4(config-if)# dialer string 8140010
R4(config-if)# dialer remote-name R1
R4(config-if)# dialer pool 1
R4(config-if)# dialer-group 1
R4(config-if)# exit
R4(config)# dialer-list 2 protocol ip list 102
```

```
R4(config)# access-list 102 permit tcp any any eq smtp
R4(config)# interface dialer 2
R4(config-if)# ip address 172.16.30.1 255.255.255.252
R4(config-if)# encapsulation ppp
R4(config-if)# ppp authentication chap
R4(config-if)# dialer string 9650001
R4(config-if)# dialer remote-name CentralOffice
R4(config-if)# dialer pool 1
R4(config-if)# dialer-group 2
R4(config-if)# exit
R4(config)# interface serial 0/0
R4(config-if)# backup interface dialer 1
```

Summary

This section summarizes the key points discussed in this lesson.

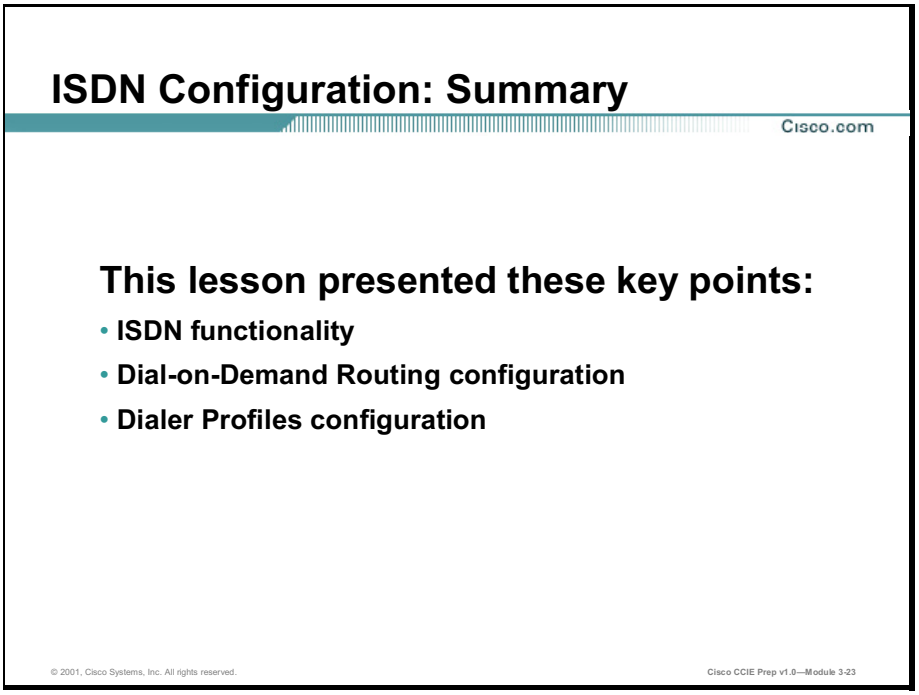

Next Steps

After completing this lesson, go to:

- PPP Features

References

For additional information, refer to these resources:

- Building Cisco Remote Access Networks (BCRAN) – Chapter 7
- DDR Configuration
  http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121cgcr/dialtsc/dtsprt5/index.htm

Lesson Assessment (Quiz)

Q1) What is the default encapsulation type on an ISDN BRI interface?

A) PPP

B) HDLC

C) ARPA

D) DDR

Q2) Which of the following is an optional component of a dialer profile?

A) Dialer interfaces

B) Dialer pool

C) Physical interfaces

D) Dialer map-class

Q3) If access-list 101 is used to specify interesting traffic, which of the following will bring up a DDR link?

```
R4(config)# access-list 101 deny eigrp any any
R4(config)# access-list 101 deny udp any any eq 520
R4(config)# access-list 101 deny tcp any any eq 21
R4(config)# access-list 101 permit ip any any
```

A) RIP

B) FTP

C) EIGRP

D) BGP

Q4) Which commands should be used on the hub for IP address negotiation? (Pick two)

A) Router(config-if)# ip address negotiated

B) Router(config)# ip local pool default

C) Router(config)# ip address-pool local

D) Router(config-if)# ip unnumbered

Q5) Which command is not needed on the physical BRI interface configuration when using dialer profiles?

A) **no ip address**

B) **encapsulation ppp**

C) **dialer pool-member**

D) **dialer-group**

PPP Features

Overview

Point-to-Point Protocol (PPP) is versatile and can be applied in a variety of situations. This lesson will examine some of the advanced features that are available when running PPP. Those features include: Password Authentication Protocol (PAP) and Challenge-Handshake Authentication Protocol (CHAP) Authentication, Multilink PPP, and PPP Callback.

Importance

Knowledge of CHAP, PAP, Multilink PPP and PPP callback are required for the Cisco Certified Internetwork Expert (CCIE) exam.

Objectives

Upon completing this lesson, you will be able to configure:

- CHAP and PAP authentication
- Multilink PPP
- PPP Callback
- Caller Identification

Learner Skills and Knowledge

To fully benefit from this lesson, you must have these prerequisite skills and knowledge:

- Passed the CCIE written qualification exam
- Completed the Building Cisco Remote Access Networks (BCRAN) course or have the equivalent knowledge

Outline

This lesson includes these sections:

- Overview
- PAP
- CHAP
- PPP Multilink
- PPP Callback
- Caller Identification
- Summary
- Lesson Assessment (Quiz)

PAP

Password Authentication Protocol (PAP) authentication can occur in bi-directional and uni-directional configurations. Each is appropriate for different scenarios.

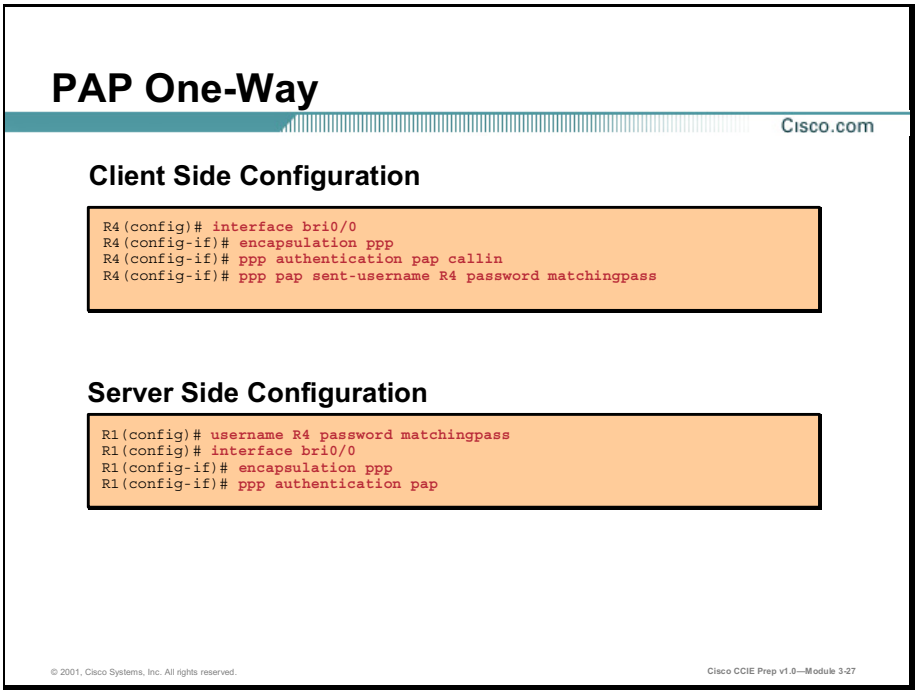

Look at a basic PAP one-way authentication configuration. Assume that you only want R1 to authenticate R4. R4 will not authenticate R1. In this scenario, you can think of R4 as the client (caller) and R1 as the server (receiver).

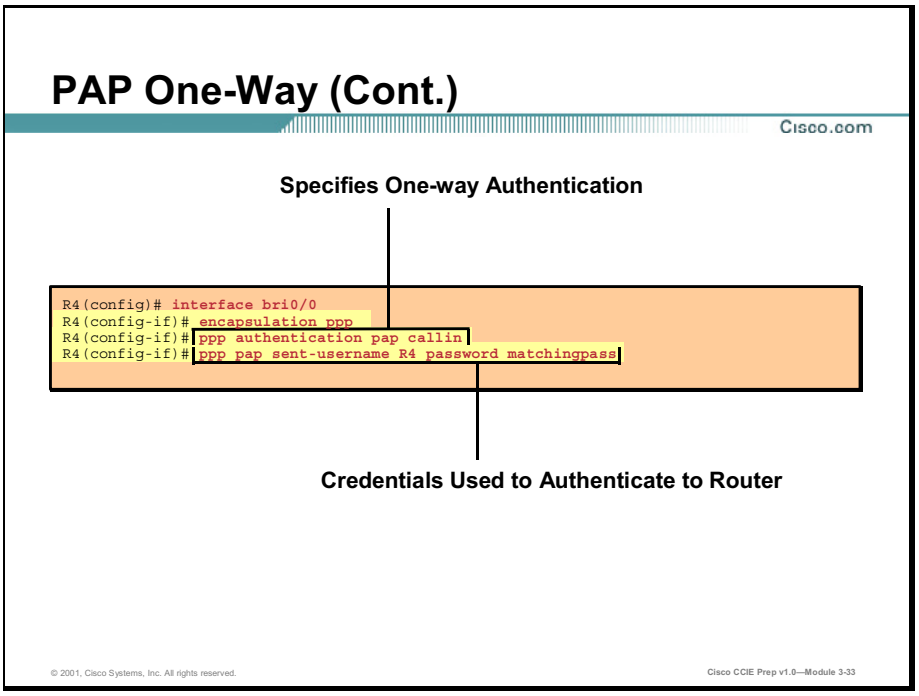

Examine the configuration for R4 (the client). This router needs to produce identification in order to gain access to resources beyond R1 (the server).

PAP requires Point-to-Point Protocol (PPP) encapsulation, which is specified first. Next, issue the command **ppp authentication pap callin**, which specifies PAP as the authentication method. The **callin** keyword specifies a one-way authentication scenario, which means R4 (the client) will not request that R1 (the server) authenticate itself.

Finally, the credentials R4 will use to authenticate to R1 are supplied. This is accomplished with the command **ppp pap sent-username R4 password matchingpass**. This statement permits outbound authentication from this client, by sending a PAP AUTH-REQ packet to R1 with the username R4 and the password matchingpass. Remember, the server (R1) must have this exact username/password in its local database in order for authentication to succeed.

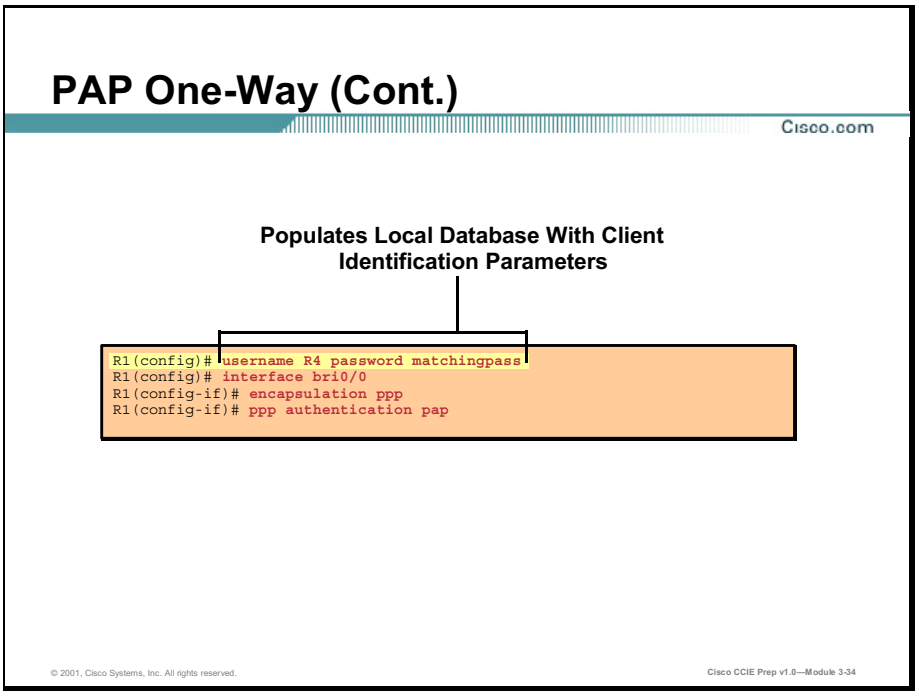

Now, take a look at R1's (the server's) configuration. First, populate the local database with the identification parameters used by the client (R4). This is performed using the command **username R4 password matchingpass**. It is important to note that you could have chosen any username/password combination. The only stipulation is they must match on the client and server.

Next, enable PPP encapsulation and specify that PAP will be the authentication method. Since R1 (the server) will not authenticate itself to any client, no further configuration is needed.

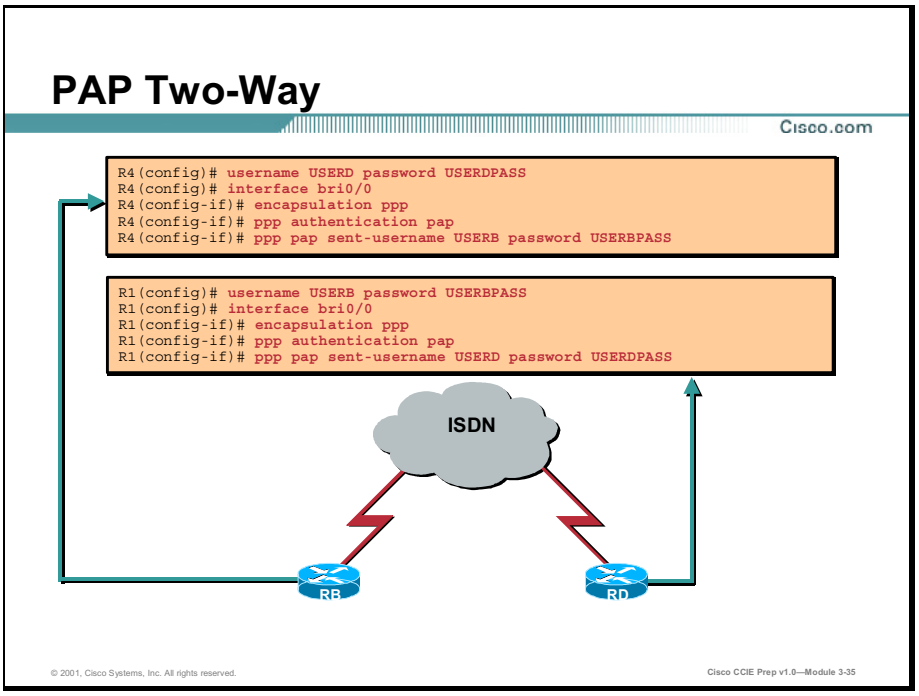

Next, take a look at two-way PAP authentication. In this example, both R4 and R1 will perform both the client and server functions. They will each provide identification (client) and request identification (server) for mutual authentication.

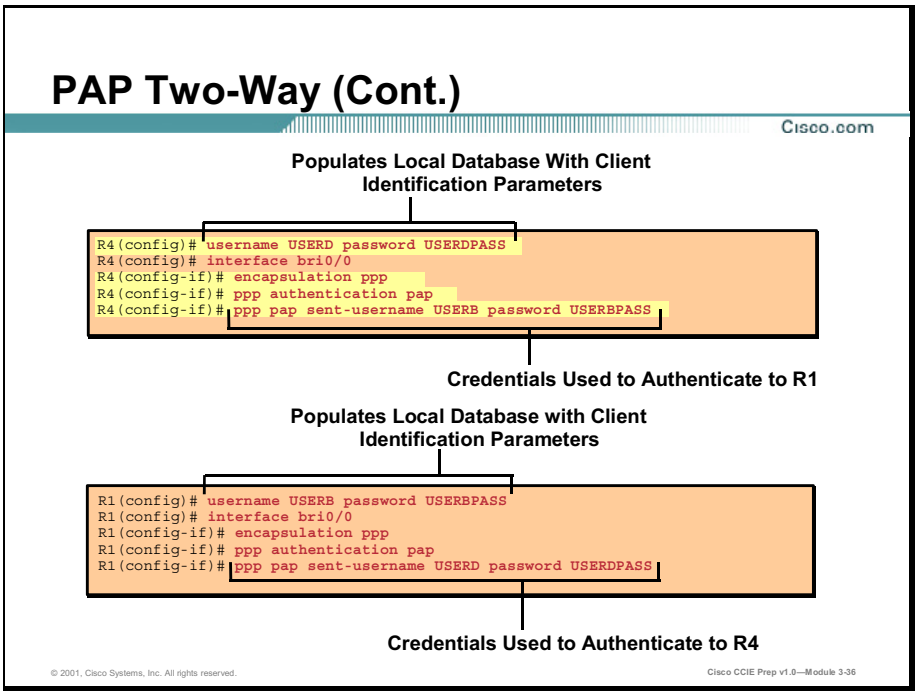

Looking at R4's configuration, you can see the local database has been populated with parameters R1 will supply as its identification. This is for the server portion of its configuration.

Next, enable PPP encapsulation, then specify the PAP authentication requirement with the command **ppp authentication pap**. Notice the keyword **callin** has been removed. The keyword **callin** is only used for one-way authentication.

Finally, R4 is supplied with the credentials it will use to authenticate to R1. This is accomplished with the command **ppp pap sent-username USERB password USERBPASS**. This statement permits outbound authentication from this client, by sending a PAP AUTH-REQ packet to R1 with the username USERB and the password USERBPASS. This was for the client portion of its configuration.

As you can see, you perform the same configuration on R1, using the correct username and password for its client/server configurations.

CHAP

Challenge-Handshake Authentication Protocol (CHAP) authentication is substantially more secure than PAP because of increased sophistication. Just like PAP, it can be configured in either a uni-directional or bi-directional setup.

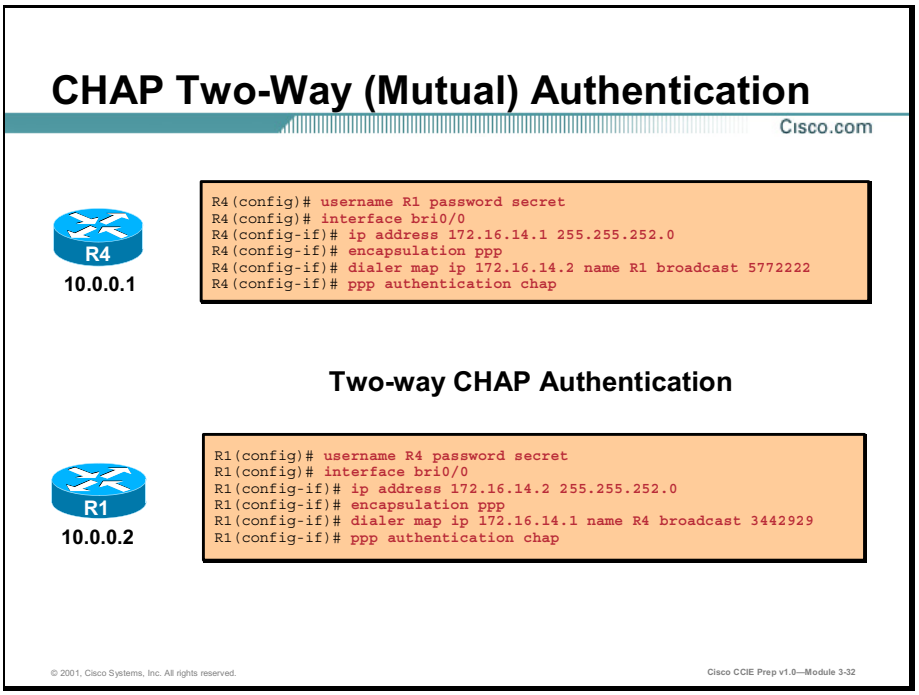

PPP negotiation involves several steps, such as Link Control Protocol (LCP) negotiation, Authentication, and Network Control Protocol (NCP) negotiation. If the two sides cannot agree on the correct parameters, then the connection is terminated. Once the link is established, the two sides authenticate each other using the authentication protocol decided on during LCP negotiation. Authentication must be successful prior to starting NCP negotiation. Shown here is a configuration showing only the relevant parameters for CHAP two-way authentication.

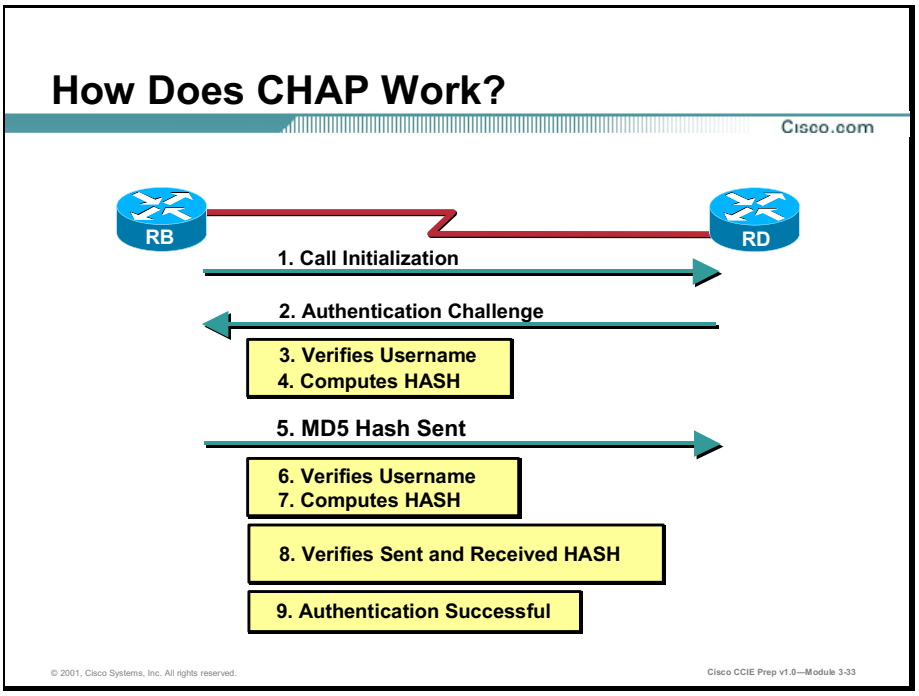

This scenario examines exactly how CHAP authentication works from the perspective of R4, who will initiate a call to R1. *

1. When R1 receives the call, it challenges R4 for authentication. By default, the hostname of the router is used to identify itself. If the **ppp chap hostname** *name* command is configured, a router uses this *name* in place of its hostname to identify itself. In this example, the challenge is labeled as it is coming from "R1."
2. R4 receives R1's challenge and looks in its local database for username "R1."
3. R4 finds an entry for “R1" and checks for a password, which is "secret." R4 uses this password and the challenge information from R1 as input parameters for the Message Digest Version 5 (MD5) hash function. The hash value is generated based on this information.
4. R4 sends the hash output value to R1.
5. R1 receives the reply and looks for the "R4" username in its local database for the password.
6. R1 finds that the password for "R4" is "secret." R1 uses the password and the challenge information sent out earlier to R4 as input parameters for the MD5 hash function. The hash function generates a hash value.
7. R1 compares the hash value it generated and the one it receives from R4.
8. Since the input parameters (challenge and password) are identical, the hash value is the same resulting in a successful authentication.

* This scenario only displayed a one-way authentication. Normally, in step 1, when R4 initiates a call it would send a challenge to R1 and the steps would be followed identically as shown.

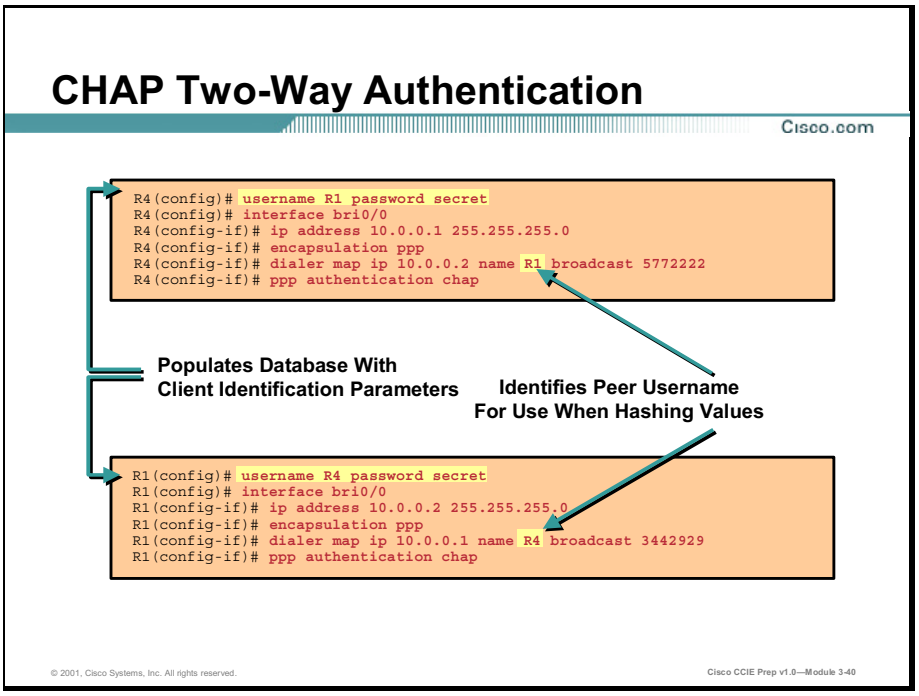

First, establish the routers' identities as well as the shared secret password. R4 populates its local user database with R1's hostname and shared secret password. R1 does the same for R4. The dialer maps also need to identify the peer's username. This is the name that will be looked up in the local database when hashing is performed to match values.

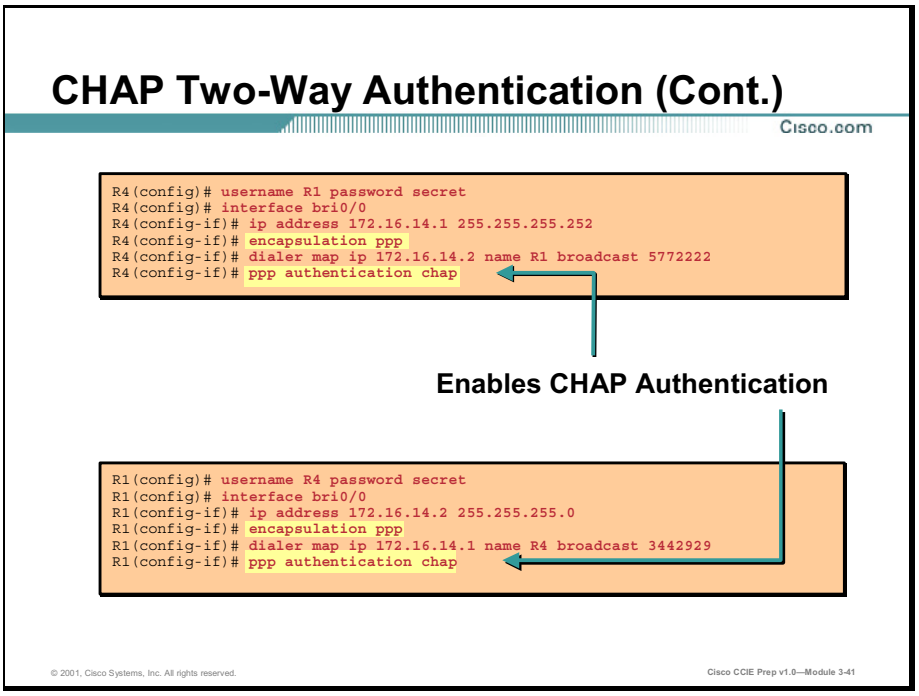

To perform mutual authentication, both routers are required to use PPP encapsulation, and then issue the **ppp authentication chap** command. This command states that you want to perform CHAP authentication as well as challenge any peer who wishes to communicate to or through the router.

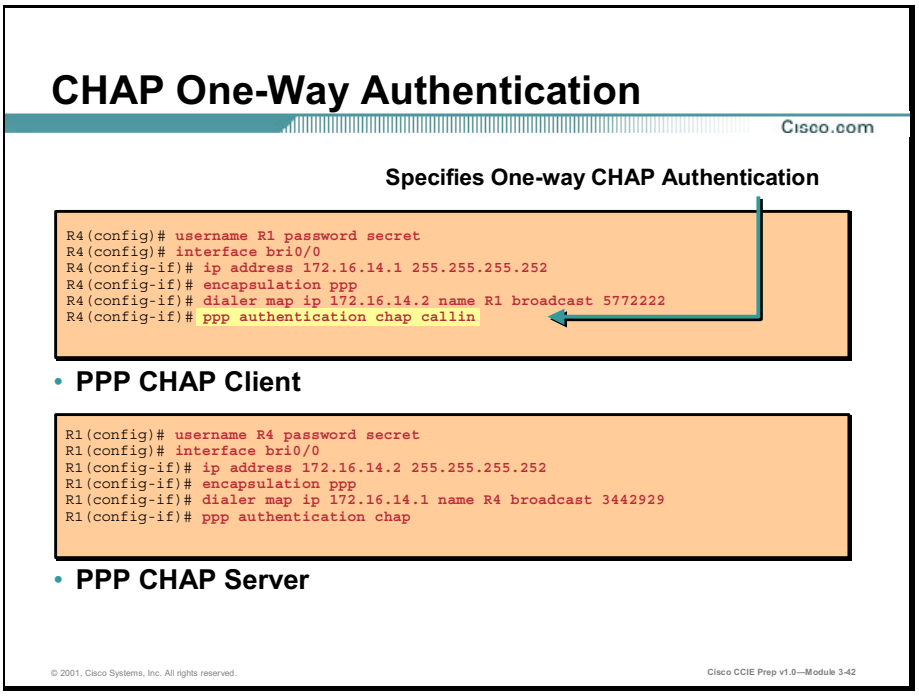

When two devices normally use CHAP authentication, each side sends out a challenge, to which the other side responds and is authenticated by the challenger. Each side authenticates one another independently. There are times when mutual authentication cannot be performed, such as the case when the initiator does not support authentication or the server does not need to authenticate to the client. In that case, you must perform one-way CHAP authentication. With one-way CHAP, the client (initiator) is authenticated, but not the server (receiver). Consider the example configuration shown, where R4 is the client and will initiate calls to R1, which is the server.

When using the **ppp authentication** command with the **callin** keyword, the Access Server (R1) will only authenticate the remote device if the remote device initiated the call (for example, if R4 "called in"). In this case, from R1's perspective, authentication is required on incoming (received) calls only. In other words, when R4 initiates a call to R1, it will not send a CHAP challenge to R1. R1 will challenge R4 (the client) to authenticate itself.

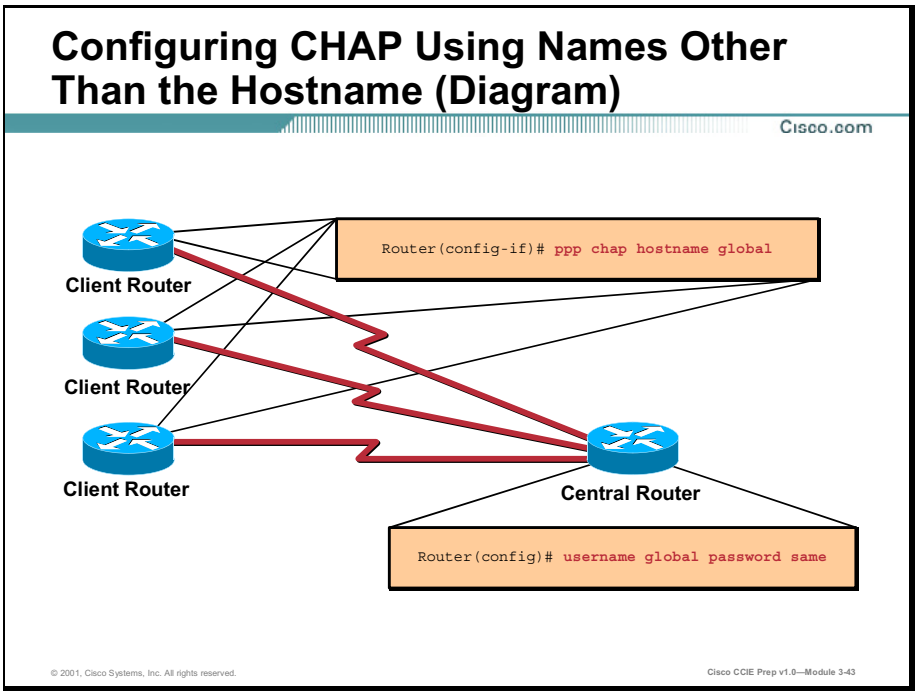

When a remote Cisco router connects to either a Cisco or a non-Cisco central router of a different administrative control, an Internet Service Provider (ISP), or a rotary of central routers, it may be necessary to configure an authentication username that is different from the hostname. In this situation, the hostname of the router is not provided or is different at different times (rotary). Also, the username that is allocated by the ISP may not be the remote router's hostname. In such a situation, the **ppp chap hostname** command is used to specify an alternate username that will be used for authentication.

For example, consider a situation where multiple remote devices are dialing into a central site. Using normal CHAP authentication, the username (which would be the hostname) of each remote device and a shared secret must be configured on the central router. In this scenario, the configuration of the central router can become lengthy and cumbersome to manage; however, if the remote devices use a username that is different from their hostname this can be avoided. The central site can be configured with a single username and shared secret that can be used to authenticate multiple dialin clients.

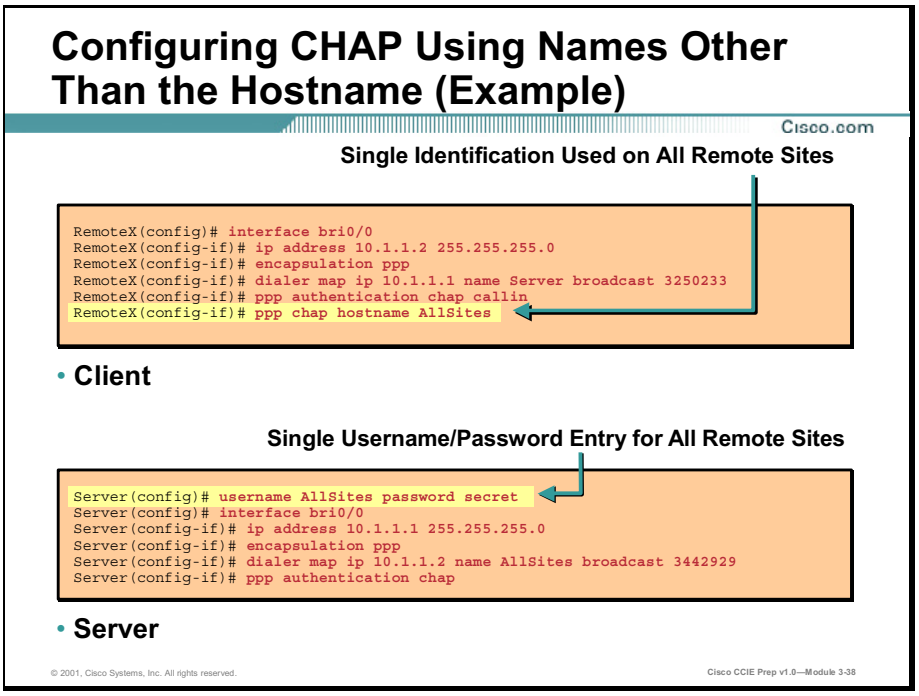

In this configuration, all remote sites will have a single identification they use to authenticate to the Access Server at the ISP. All remote sites use the username of “AllSites” with the shared secret password of “secret”.

PPP Multilink

Multilink PPP (MPPP) provides a method for spreading traffic across multiple physical Wide Area Network (WAN) links while providing packet fragmentation and reassembly, proper sequencing, multivendor interoperability, and load balancing on inbound and outbound traffic.

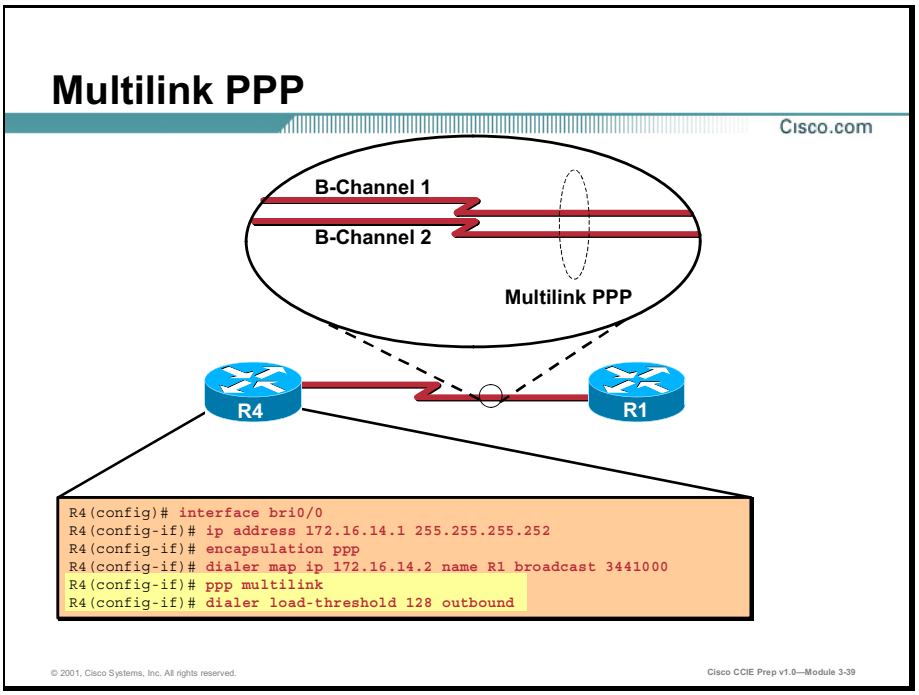

There are two basic methods for configuring MPPP. The first method, which is the simpler of the two, can be used only on Integrated Services Digital Network (ISDN) Basic Rate Interfaces (BRIs) and Primary Rate Interfaces (PRIs). The second involves the creation of a Dialer interface, which can be used with any type of WAN connection. The example shows an MPPP configuration using the first method.

The two commands highlighted work in concert to provide all the features of MPPP.

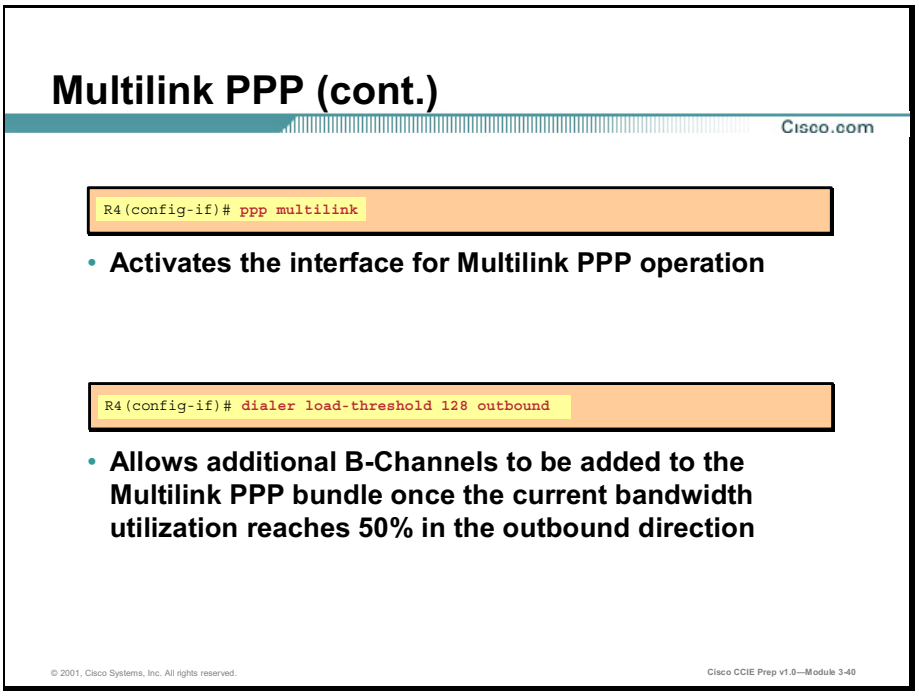

The command **ppp multilink** activates the interface for MPPP operation and allows negotiation of Multilink PPP at connect time, thus establishing a single-channel MPPP bundle. However, this command alone is not sufficient enough to take advantage of the fragmentation, load balancing, or bandwidth-on-demand features of the Multilink PPP.

The **dialer load-threshold** ***load*** command sets the point at which additional B channels will be added to the MPPP bundle. When the total load of all "up" B channels (*n*) is greater than the load threshold, the dialer interface (in this case, BRI 0/0) adds an extra channel to the multilink bundle. In a similar way, if the total load for all the "up" B channels minus one (*n* - 1) is at or below the threshold, the additional channels will be taken back down.

The ***load*** argument is the average load for the interface; it is a value from 1 (unloaded) to 255 (fully loaded). As shown in the above example a 50% load threshold is achieved by configuring the "load" argument with a value of 128. The load argument is expressed as a percentage n/255 where n is the value configured.

The **outbound** argument sets the load calculation to be made only on outbound traffic. The **inbound** argument does the same, but for inbound traffic only. Using the **either** argument sets the load as the larger of the outbound and inbound loads.

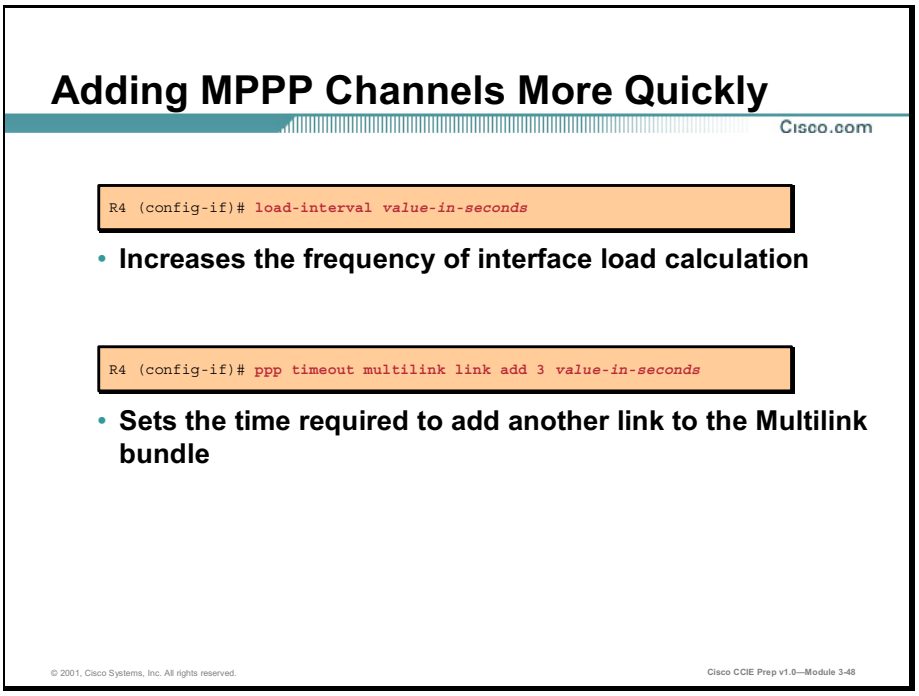

On a Primary Rate Interface (PRI) ISDN circuit, it is often desirable to add additional B channels as quickly as possible. If Multilink Point-to-Point Protocol (MPPP) is configured on the PRI interface, the following commands allow you to add B channels to the Multilink bundle more quickly:

```
R4(config-if)# load-interval 30
```

This command increases the frequency of the interface load calculation. By default, the interface load is calculated as an exponential average over the last five minutes. By setting the "load-interval" parameter to 30 seconds, you force a more frequent calculation of the interface load. This results in an earlier detection of changes in the interface load. By shortening the length of time during which you compute the interface load, you also shorten the time required to bring up additional B channels because the router will detect an increase of the interface load sooner.

```
R4(config-if)# ppp timeout multilink link add 3
```

This command sets the time required to add another link to the multilink bundle. This timer determines how long PPP multilink waits after adding a link to the bundle before adding additional links due to the load threshold being exceeded. The minimum possible value is 1 second. The default is 30 seconds.

Note These commands are also useful for testing Multilink PPP on BRI interfaces. These commands will reduce the time required to see if Multilink PPP is working properly. Time-saving tips like these can be extremely helpful in the CCIE Lab.

PPP Callback

PPP callback provides a client-server relationship between the end points of a point-to-point connection. PPP callback allows a router to request that a peer router call back.

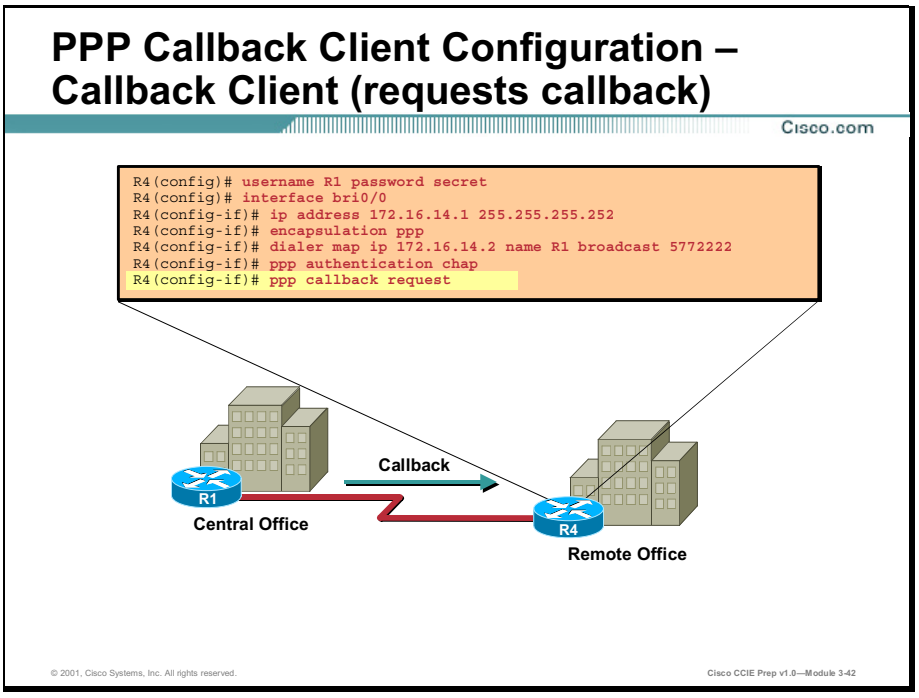

Client configuration is very simple. You request the server to call you back using the command **ppp callback request**, as shown in the example above.

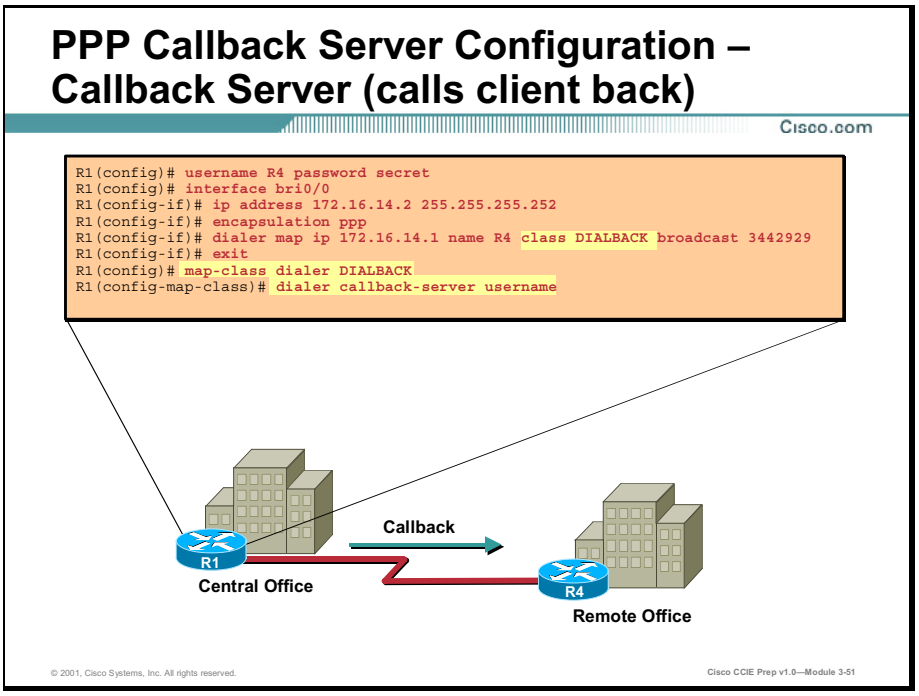

Server configuration is a little more complex. First, configure the server to accept callback requests with the **ppp callback accept** command.

The command **dialer callback-secure** is to perform a couple of key security functions.

- Disconnect calls that are not properly configured for callback
- Disconnect any unauthenticated dial-in users

It is not a required component for callback to succeed, but is highly recommended in all callback configurations.

To enable an interface to make return calls when callback is successfully negotiated, issue the **dialer callback-server** command via a map class. A map class is used to define a template of configuration parameters for PPP callback. The keyword **username** identifies the return call by looking up the authenticated host name in the **dialer map** command. In this case, it would be R4.

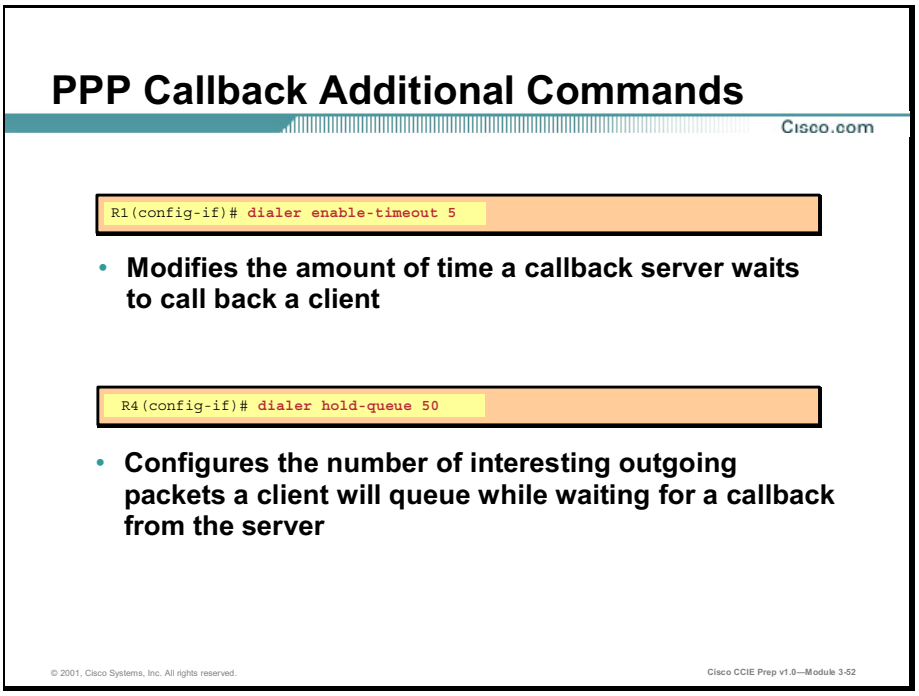

By default, the callback server will wait 15 seconds before it attempts to initiate a callback to the client. You can modify this timer with the **dialer enable-timeout** command. To set it to 5 seconds issue the command:

```
R1(config-if)# dialer enable-timeout 5
```

During these few seconds, the packets being sent by the client will be dropped. You can have the client queue these packets until the link is established, then send them using the command below.

To allow *interesting* outgoing packets to be queued until a connection is established, use the **dialer hold-queue** command in interface configuration mode.

```
R4(config-if)# dialer hold-queue 50
```

Caller Identification

Caller ID screening allows the initial incoming call from the client to the server to be accepted or rejected based on the caller ID message contained in the ISDN setup message. Caller ID screening also allows the server to initiate a callback to the calling client.

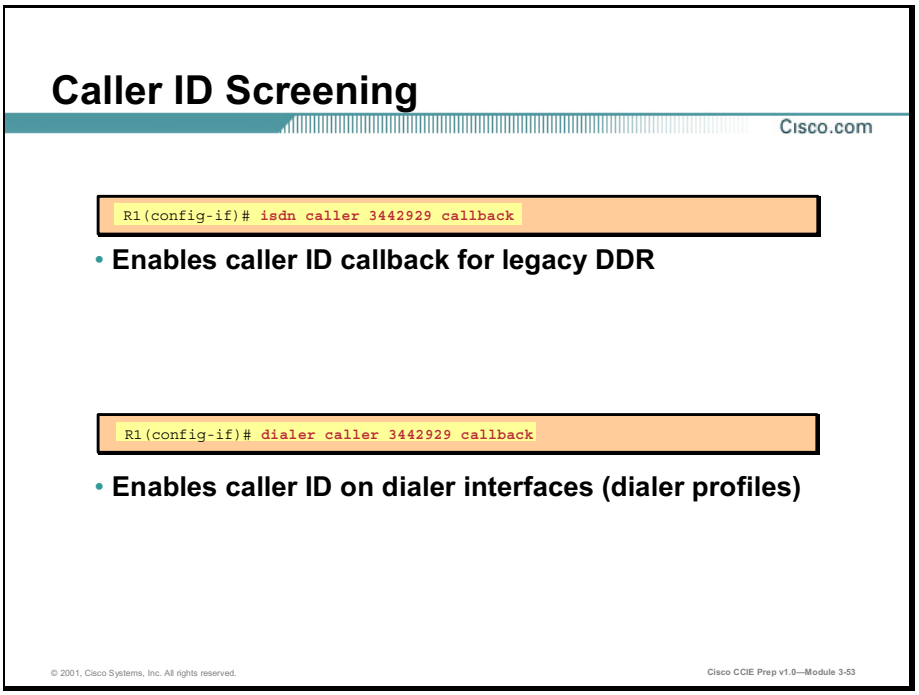

To configure caller ID screening and optionally enable ISDN caller ID callback for legacy DDR, use the **isdn caller** interface configuration command.

```
R1(config-if)# isdn caller 3442929 callback
```

To configure caller ID screening and optionally enable ISDN caller ID callback for dialer profiles, use the **dialer caller** interface configuration command.

```
R1(config-if)# dialer caller 3442929 callback
```

Summary

This section summarizes the key points discussed in this lesson.

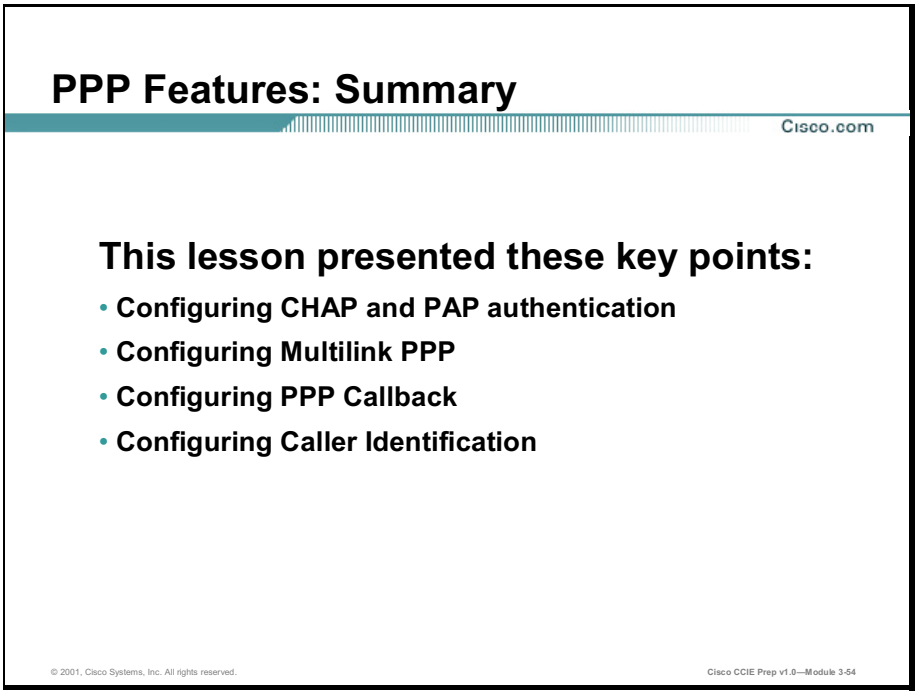

Next Steps

After completing this lesson, go to:

- Using ISDN as a Backup Connection

References

For additional information, refer to these resources:

- http://www.cisco.com/warp/public/131/ppp_callin_hostname.html
- http://www.cisco.com/warp/public/779/smbiz/service/configs/isdn/isdn_configs.htm

Lesson Assessment (Quiz)

Q1) Which authentication method sends a clear-text password?

A) CHAP

B) PAP

C) PPP

D) MPPP

Q2) What authentication mechanism should be used if the destination device supports encrypted hashed messages, but cannot initiate authentication?

A) PAP one-way

B) PAP two-way

C) CHAP one-way

D) CHAP two-way

Q3) Which command changes how frequently MPPP calculates the need for additional B channels?

A) ppp timeout multilink link add

B) ppp multilink

C) load-interval

D) dialer load-threshold

Q4) The “sent-username” feature is used with which two authentication schemes?

A) PAP one-way

B) PAP two-way

C) CHAP one-way

D) CHAP two-way

Q5) What CHAP command should be used on a hub router that requires a different hostname be sent to remote sites?

A) ppp chap altname

B) ppp authentication chap no username

C) ppp chap hostname

D) ppp chap sent-username

Using ISDN as a Backup Connection

Overview

Integrated Services Digital Network (ISDN) offers high-bandwidth, inexpensive backup media to higher bandwidth lines, such as T1. This lesson will examine the Backup Interface and Dialer Watch features, along with suggestions for implementing each one.

Importance

This lesson details the complex range of uses and implementations for ISDN dial backup configurations.

Objectives

Upon completing this lesson, you will be able to:

- Configure Backup Interfaces
- Configure Dialer Watch
- Describe when to use a given dial backup implementation

Learner Skills and Knowledge

To fully benefit from this lesson, you must have these prerequisite skills and knowledge:

- Passed the Cisco Certified Internetwork Expert (CCIE) written qualification exam
- Completed the Building Cisco Remote Access Networks (BCRAN) course or have the equivalent knowledge

Outline

This lesson includes these sections:

- Overview
- Floating Static Routes
- Backup Interface
- Backup Load
- Dialer Watch
- Dialer Watch Operation
- Dialer Watch Configuration
- Characteristics of the Backup Methods
- Summary
- Lesson Assessment (Quiz)

Floating Static Routes

Floating static routes are an enhancement to static routes that use administrative distance to appropriately weight the backup route in relation to routes learned through dynamic routing protocols.

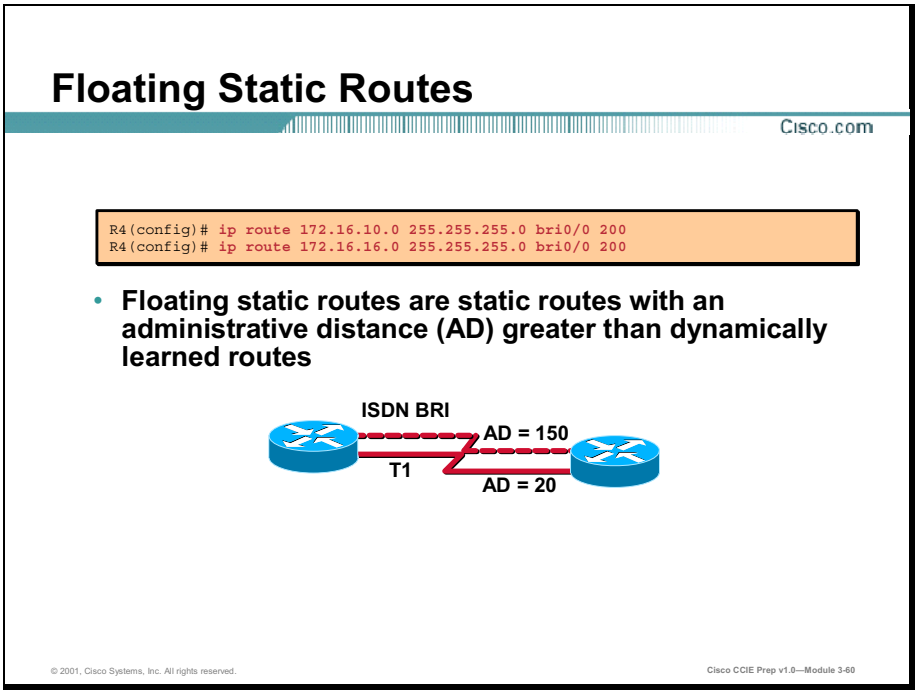

When ISDN is implemented in the real world, it is usually with the use of floating static routes. Floating static routes are static routes that have an Administrative Distance (AD) greater than the administrative distance of dynamically learned routes. An Administrative Distance can be assigned to a static route so that the static route is less desirable than a dynamic route. In this manner, the static route is not used when the dynamic route is available. However, if the dynamic route is lost, the static route can take over, and traffic can be sent through this alternate route. If this alternate route is provided using a Dial-on-Demand Routing (DDR) interface, then the DDR interface can be used as a backup mechanism.

Implementing floating static routes is quick, simple, and easy to test. If your primary link is routing data for the 172.16.10.0/24 and 172.16.16.0/24 networks, you can implement floating static routes as follows:

```
R4(config)# ip route 10.20.20.0 255.255.255.0 bri0/0 200
R4(config)# ip route 10.30.30.0 255.255.255.0 bri0/0 200
```

If R4 ever loses its dynamically learned routes, which should have an administrative distance less than 200, the floating static routes will come into effect and route traffic over the bri0/0 circuit. If and when the dynamically learned routes enter the routing table again, due to their lower AD, they will be the preferred entries once again.

Backup Interface

A backup interface is an interface that stays idle until certain circumstances occur; then it is activated.

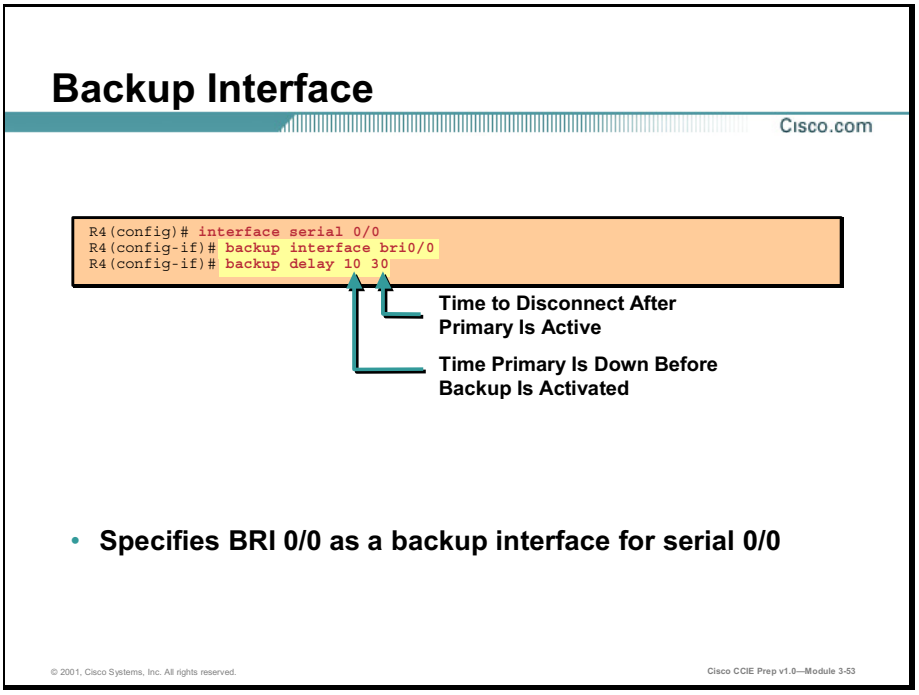

This example shows R4's Basic Rate Interface (BRI) interface being used to backup its primary link, which is the serial 0/0 interface

The **backup interface** command is placed under the primary link. This is the link that needs to be backed up in case of failure. Here you are specifying that the interface backing up the primary link is bri0/0.

Next, specify how quickly the backup interface would be activated upon failure of the primary interface. In the scenario, the **backup delay** is configured for 10 seconds. The backup interface will come up 10 seconds after it notices the primary link has failed. The backup interface is also configured to disconnect 30 seconds after the primary link is again operational.

Testing the backup interface is a little more difficult. A simple shutdown of the interface will not cause the IOS software to see this as a link failure. To test the backup interface, you must physically remove the cable from the serial interface. This is not the case if the **backup interface** command is configured on a Frame Relay point-to-point subinterface. If the **shutdown** command is performed on the Frame Relay interface of the router on the other side of the Permanent Virtual Circuit (PVC), it will cause the PVC to go "Inactive", causing the point-to-point subinterface to go down/down. This state will trigger the backup interface.

Backup Delay

The **backup delay** command can be used to control how quickly a secondary line is brought up.

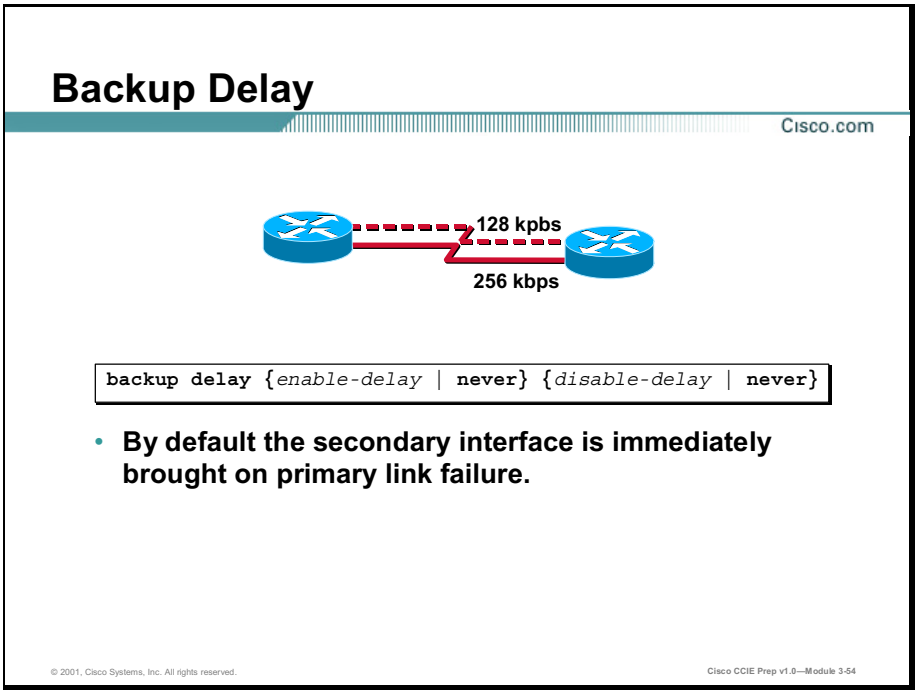

You can configure a value that defines how much time should elapse before a secondary line status changes after a primary line status change. This means that you can define two delays:

- A delay that specifies the amount of time after the primary line goes *down*, but before the secondary line is activated
- A delay that specifies the amount of time after the primary line comes *up*, but before the secondary line is deactivated

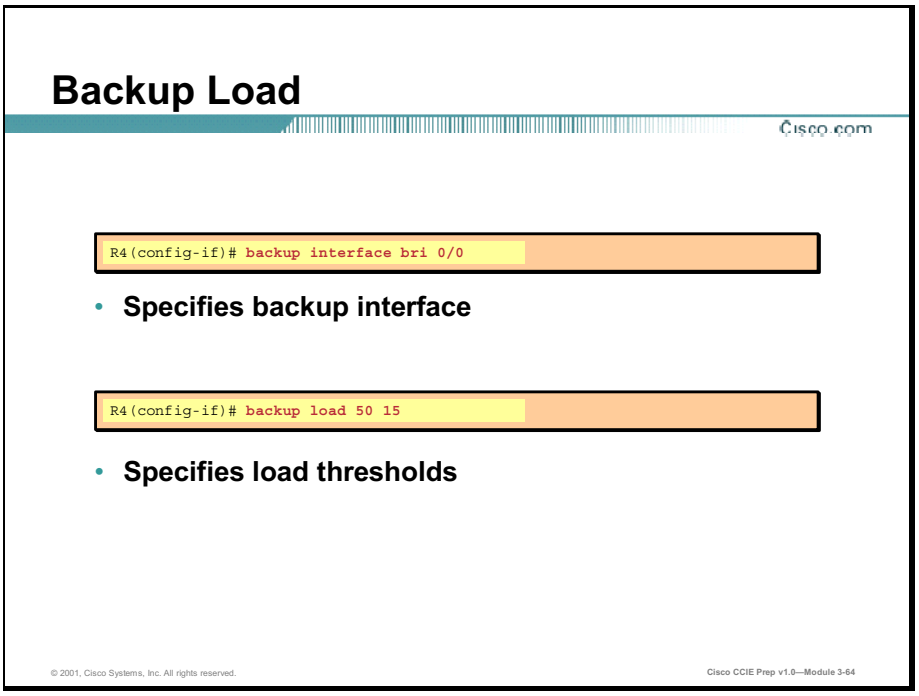

The command shown here specifies that the BRI interface will not backup the serial interface upon failure, but instead provide additional bandwidth when certain load thresholds are met.

In this example, the BRI interface will activate when 50 percent of the available bandwidth on the serial interface is reached. It is very important that the actual bandwidth of the serial 0/0 interface be set with the **bandwidth** command, otherwise the enable threshold might never be met and Bandwidth-On-Demand (BOD) will not occur. The BRI interface is also configured to deactivate when the available bandwidth on the serial interface drops below 15 percent.

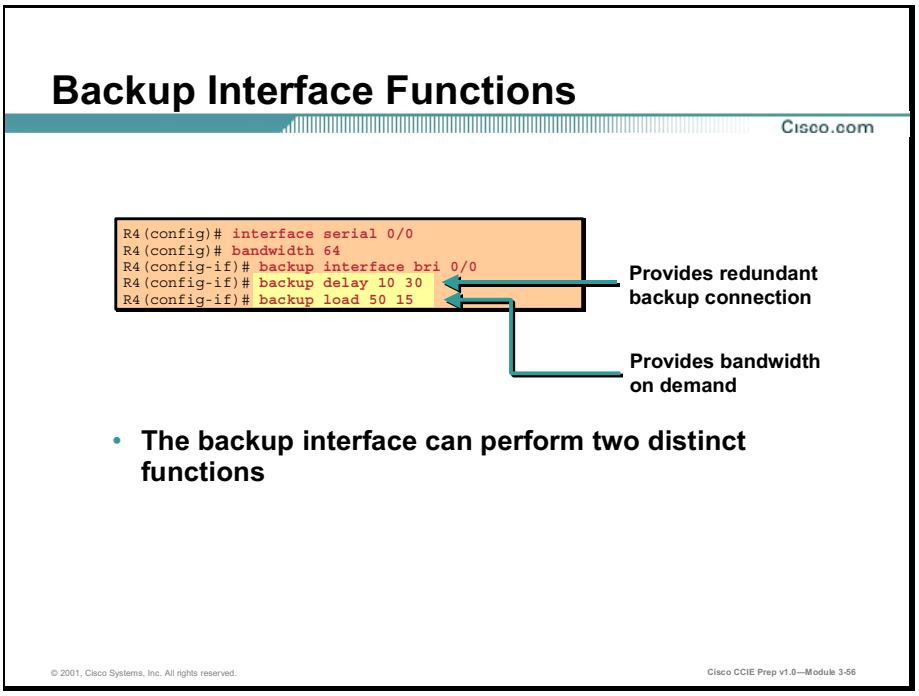

The backup interface can be used to perform two distinct functions:

- Backup upon failure - Using the **backup delay** command. The backup interface will not be activated until the primary link fails.
- Bandwidth on Demand - Using the **backup load** command. The backup interface will not be activated until the primary link bandwidth reaches a certain load threshold.

Note	You can configure a secondary line to be both backup upon failure and bandwidth on demand at the same time to take advantage of both functions.

Dialer Watch Configuration

Dialer watch configuration is built by having an interface monitor a specified route or set of routes.

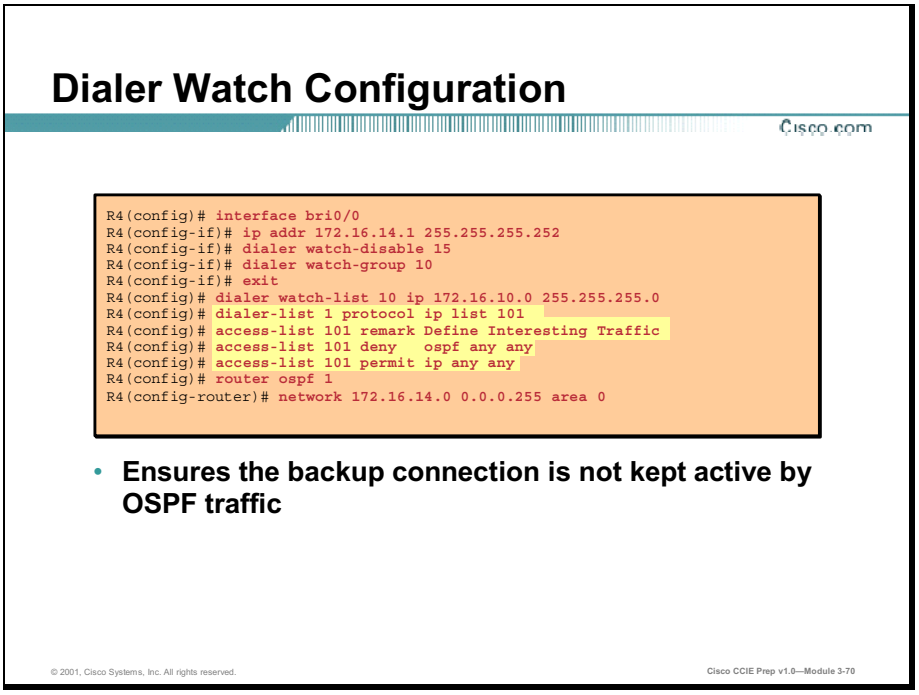

In this example, the BRI0/0 interface is part of Open Shortest Path First (OSPF) area 0. You are learning other OSPF routes over the primary link (not shown). You have network 172.16.10.0/24 in the routing table, which is going to be a watched route. Since you are using OSPF, you do not want the OSPF multicast hello packets to constantly bring up the ISDN line, so you mark OSPF traffic as uninteresting in the dialer list.

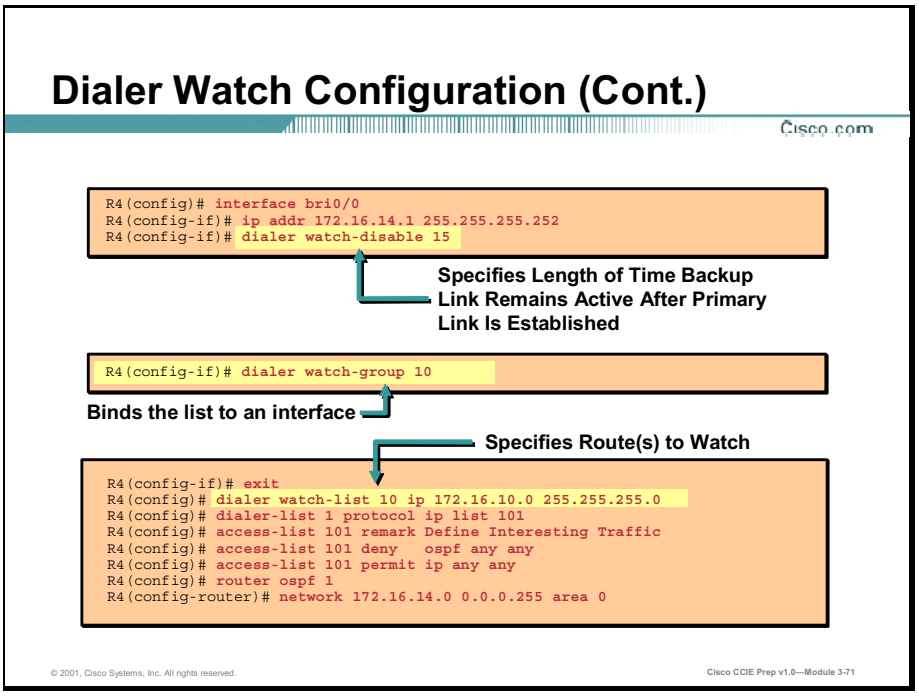

The **dialer watch-list** specifies the route to monitor, 172.16.10.0/24 in this case. You apply that list to the BRI interface with the **dialer watch-group** command. The **dialer watch-disable** command delays disconnecting the backup interface for 15 seconds after the primary interface comes back up.

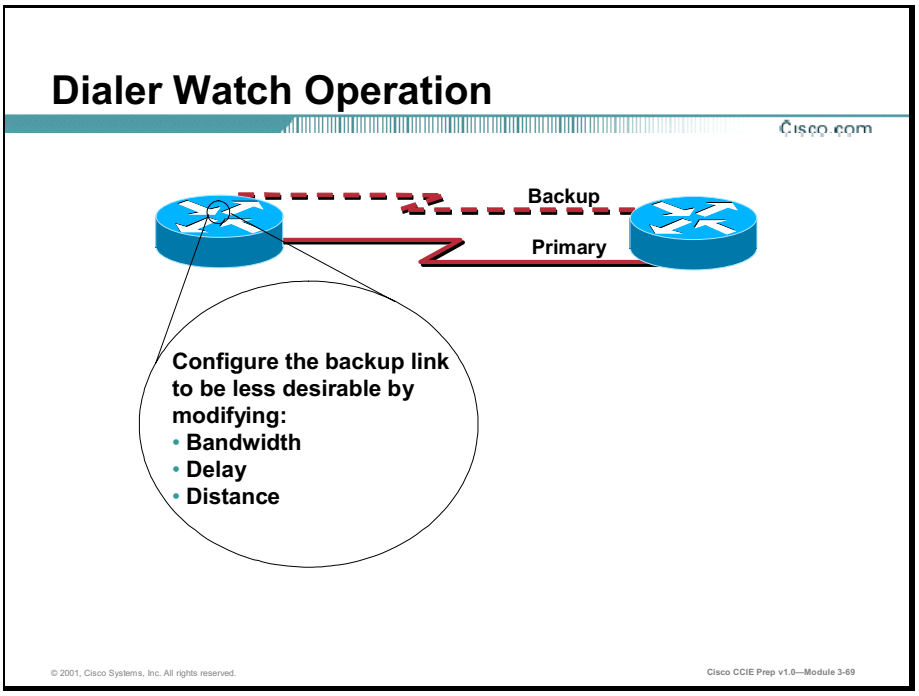

Configure the backup link to have a lower routing preference than the primary link. This is done because when the primary link becomes available again, the dynamic routing protocol should prefer the primary over the backup link and not attempt to load-balance across the two links, thus keeping the backup link up indefinitely. The backup link can be configured to be less preferable with any of the following commands; **bandwidth**, **delay** or **distance** as appropriate.

Characteristics of the Backup Methods

This section covers three methods for configuring backup interfaces in a DDR topology.

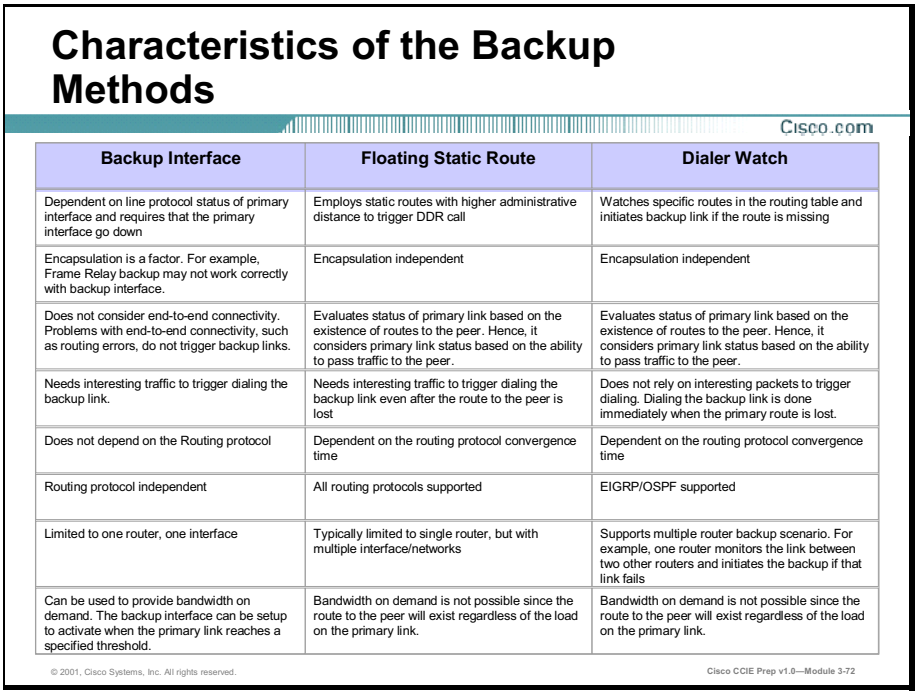

This table summarizes the characteristics of the three backup methods. You can use it to compare and evaluate the appropriate backup method to use in a certain situation.

Table 3-1: Backup Methods

Backup Interface	Floating Static Route	Dialer Watch
Dependent on line protocol status of primary interface and requires that the primary interface go down	Employs static routes with higher administrative distance to trigger DDR call when dynamic routes are lost	Watches specific routes in the routing table and initiates backup link if the route is missing
Encapsulation is a factor. For example, Frame Relay backup may not work correctly with backup interface.	Encapsulation independent	Encapsulation independent
■ Does not consider end-to-end connectivity ■ Problems with end-to-end connectivity, such as routing errors, do not trigger backup links	■ Evaluates status of primary link based on the existence of routes in the routing table ■ It considers primary link status based on the ability to route traffic to the peer	■ Evaluates status of primary link based on the existence of routes in the routing table ■ It considers primary link status based on the ability to route traffic to the peer
Requires interesting traffic to trigger dialing the backup link	Needs interesting traffic to trigger dialing the backup link even after the route to the peer is lost	■ Does not rely on interesting packets to trigger dialing ■ Dialing the backup link is done immediately when the primary route is lost
Does not depend on the routing protocol	Dependent on the routing protocol convergence time	Dependent on the routing protocol convergence time
Routing protocol independent	All routing protocols supported	Only Enhanced Interior Gateway Routing Protocol (EIGRP) and OSPF are supported
Limited to one router, one interface	Typically limited to single router, but with multiple interface/networks	■ Supports multiple router backup scenario ■ For example, one router monitors the link between two other routers and initiates the backup if that link fails.
■ Can be used to provide Bandwidth On Demand (BOD) ■ The backup interface can be setup to activate when the primary link reaches a specified load threshold.	Bandwidth on demand is not possible since the route to the peer will exist regardless of the load on the primary link.	Bandwidth on demand is not possible since the route to the peer will exist regardless of the load on the primary link

Summary

This section summarizes the key points discussed in this lesson.

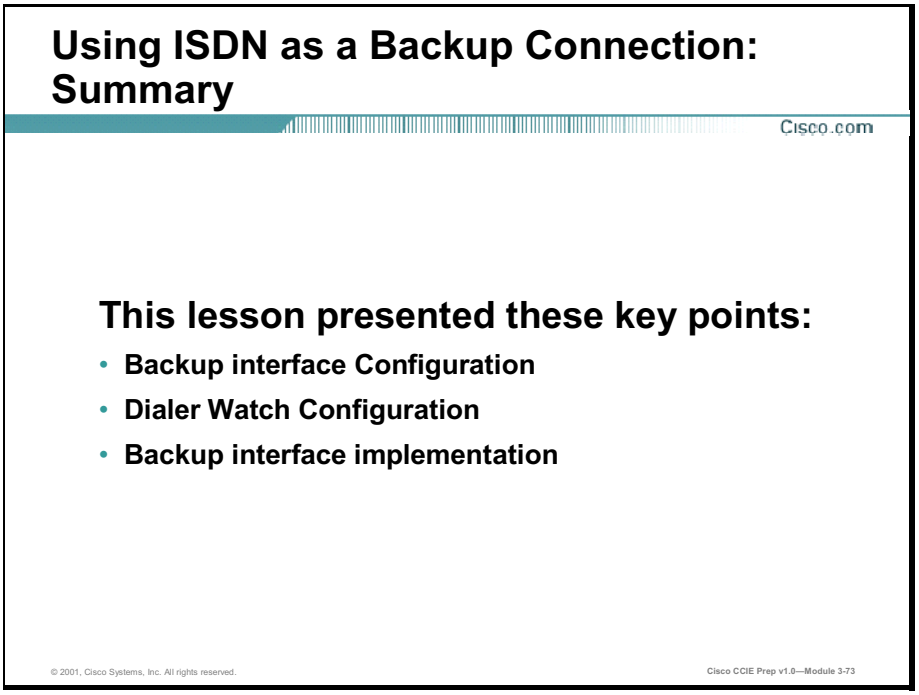

Next Steps

After completing this lesson, go to:

- Troubleshooting

References

For additional information, refer to these resources:

- http://www.cisco.com/warp/public/793/access_dial/bri_isdn_11048.html
- http://www.cisco.com/warp/public/793/access_dial/hdlc_12497.html
- http://www.cisco.com/warp/public/793/access_dial/backup_11047.html

Lesson Assessment (Quiz)

Q1) Which backup configuration method uses a static route configured with a higher administrative distance than that of a dynamically learned route to the same location?

A) Backup interface

B) Dialer watch

C) Floating static routes

D) Backup static routes

Q2) Which backup configuration monitors the status of a route within the routing table?

A) Backup interface

B) Dialer watch

C) Floating static routes

D) None of the above

Q3) Which routing protocols are supported with dialer watch?

A) RIP

B) OSPF

C) EIGRP

D) BGP

Q4) Which backup mechanism supports Bandwidth-On-Demand (BOD)?

A) Backup interface

B) Dialer watch

C) Floating static routes

D) All of the above

Q5) Which backup mechanism does not require interesting traffic to initiate a DDR call?

A) Backup interface

B) Dialer watch

C) Floating static routes

D) All of the above

Troubleshooting

Overview

Successful configuration of ISDN relies on effective troubleshooting in the event of a problem. This lesson examines techniques for troubleshooting ISDN configuration.

Importance

Show and **debug** commands are the primary tool for troubleshooting on Cisco routers.

Objectives

Upon completing this lesson, you will be able to:

- Explain the differences between specific **show** and **debug** commands
- Apply the appropriate **debug** and **show** when troubleshooting ISDN connectivity

Learner Skills and Knowledge

To fully benefit from this lesson, you must have these prerequisite skills and knowledge:

- Passed the CCIE written qualification exam
- Completed the Building Cisco Remote Access Networks (BCRAN) course or have the equivalent knowledge
- Completed the Cisco Internetwork Troubleshooting (CIT) course or have the equivalent knowledge

Outline

This lesson includes these sections:

- Overview
- **Show** Commands
- **Debug** Commands
- Summary
- Lesson Assessment (Quiz)

Show Commands

Show commands are critical for identifying the configuration and status of ISDN.

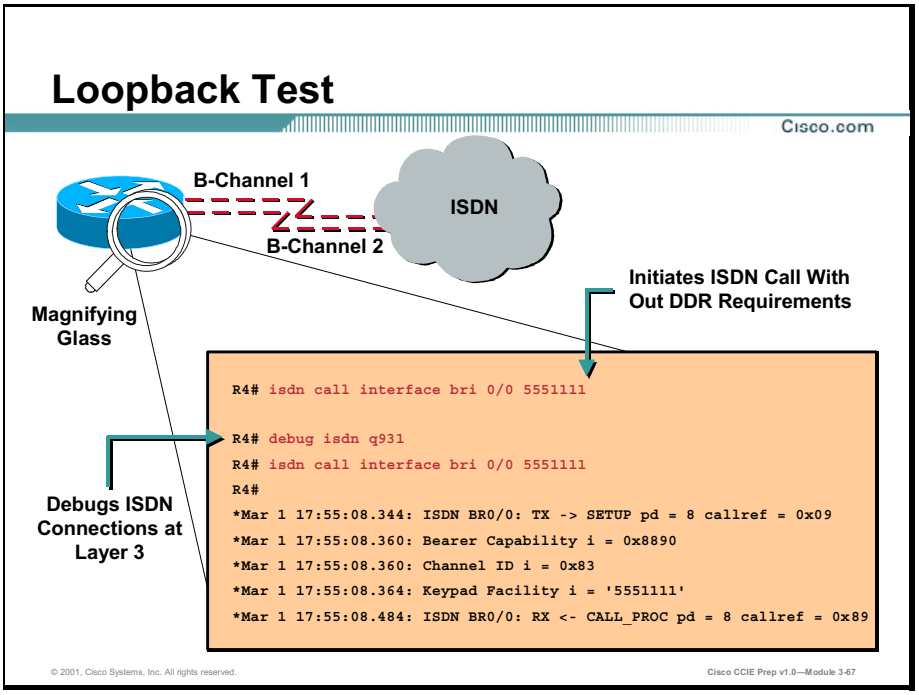

If you are experiencing problems with your BRI circuit, the first step is to perform a loopback test call.

With a loopback call, the router dials the ISDN number of its own BRI interface. The call proceeds to the telco cloud, where the telco switches the call to the second BRI channel. The router now sees this call as an incoming call on the second channel. Therefore, the router verifies that it can both send and receive ISDN calls.

A loopback call tests the ability of the router to initiate and terminate an ISDN call. A successful loopback call gives you a strong indication that the ISDN circuit to the telco cloud is functioning correctly.

The following is an annotated example of a successful loopback call. The command **isdn call** (introduced in Cisco IOS software 12.0(3)T) enables outgoing isdn calls without the DDR requirements such as interesting traffic and routes. This command can only be used for testing of the ISDN circuit and cannot be used to pass traffic or as a substitution for a proper DDR configuration. This command allows you to verify that the ISDN circuit, especially Layer 3, is functioning.

```
R4# isdn call interface bri 0/0 5551111
!--- the router will dial 5551111 (the ISDN number of the router's own BRI)
R4#
*Mar  1 17:55:08.344: ISDN BR0/0: TX ->  SETUP pd = 8  callref = 0x09
!--- Q931 Setup message is Transmitted (TX) to the telco switch
*Mar  1 17:55:08.360:         Bearer Capability i = 0x8890
*Mar  1 17:55:08.360:         Channel ID i = 0x83
```

```
*Mar  1 17:55:08.364:          Keypad Facility i = '5551111'
*Mar  1 17:55:08.484: ISDN BR0/0: RX <-  CALL_PROC pd = 8  callref = 0x89
!--- Call Proceeding message is Received (RX) from the telco switch.
!--- The switch is now processing the call.
*Mar  1 17:55:08.488:          Channel ID i = 0x89
*Mar  1 17:55:08.516: ISDN BR0/0: RX <-  SETUP pd = 8  callref = 0x12
!--- A Setup message is Received (RX) from the switch. This message is for the
!--- incoming call. Remember that the router sent a Setup message (for the
!--- outgoing call) and now receives a SETUP message for the same call
*Mar  1 17:55:08.516:          Bearer Capability i = 0x8890
*Mar  1 17:55:08.520:          Channel ID i = 0x8A
*Mar  1 17:55:08.520:          Signal i = 0x40 - Alerting on - pattern 0
*Mar  1 17:55:08.532:          Called Party Number i = 0xC1, '5551111'
*Mar  1 17:55:08.532:          Locking Shift to Codeset 5
*Mar  1 17:55:08.532:          Codeset 5 IE 0x2A  i = 0x808001038001118001, '<'
*Mar  1 17:55:08.564: ISDN BR0/0: Event: Received a DATA call from  on B2 at 64
Kb/s
*Mar  1 17:55:08.620: %DIALER-6-BIND: Interface BRI0/0:2 bound to profile
Dialer1
*Mar  1 17:55:08.652: ISDN BR0/0: TX ->  CALL_PROC pd = 8  callref = 0x92
! --- Transmit (TX) a Call Proceeding message for the incoming call
*Mar  1 17:55:08.652:          Channel ID i = 0x8A
*Mar  1 17:55:08.700: %LINK-3-UPDOWN: Interface BRI0/0:2, changed state to up
*Mar  1 17:55:08.988: ISDN BR0/0: TX ->  CONNECT pd = 8  callref = 0x92
! --- Transmit (TX) a Connect message for the incoming call
*Mar  1 17:55:08.988:          Channel ID i = 0x8A
*Mar  1 17:55:09.040: ISDN BR0/0: RX <-  CONNECT_ACK pd = 8  callref = 0x12
! --- Receive (RX) a Connect Acknowledgment for the incoming call
*Mar  1 17:55:09.040:          Channel ID i = 0x8A
*Mar  1 17:55:09.040:          Signal i = 0x4F - Alerting off
*Mar  1 17:55:09.064: ISDN BR0/0: RX <-  CONNECT pd = 8  callref = 0x89
! --- Receive (RX) a Connect for the outgoing call
*Mar  1 17:55:09.076: ISDN BR0/0: TX ->  CONNECT_ACK pd = 8  callref = 0x09
*Mar  1 17:55:09.080: %LINK-3-UPDOWN: Interface BRI0/0:1, changed state to up
*Mar  1 17:55:09.104: %DIALER-6-BIND: Interface BRI0/0:1 bound to profile BRI0/0
*Mar  1 17:55:09.112: %ISDN-6-CONNECT: Interface BRI0/0:1 is now connected to
5551111
! ---  Call is now connected. Loopback call is successful
```

Note	During the loopback call, the router performs as both the Called Router as well as the Calling Router (albeit on different B-channels). It is important that you keep track of these "dual roles" when interpreting the **debug isdn q931** output. For example, the router transmits a setup message (TX -> SETUP) and receives one as well (RX <- SETUP). The transmitted SETUP should be associated with the outgoing call while the received SETUP message is associated with the incoming call.

In the above example, you dialed the number for the first B-channel. However, the telco recognizing that the first B-channel was busy (since it was making the call), switched the call to the second B-channel and the connection was completed successfully. However, an incorrectly configured telco switch can result in a failure of the loopback call, due to the switch trying to assign the call to the first channel (which is busy making the call). The telco should correct this problem. However, as a workaround solution, specify the second B-channel number in the **isdn call** command. If the loopback call succeeds and the call to the remote end continues to fail, contact the telco for further troubleshooting assistance with your BRI circuit.

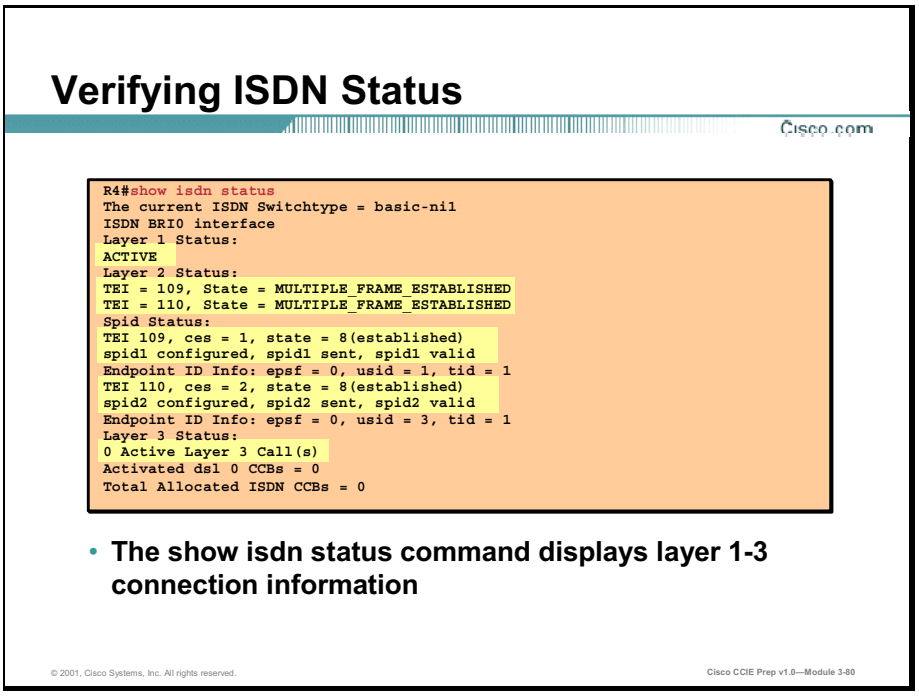

Use the **show isdn status** command to verify that ISDN BRI Layer 1 is ACTIVE, LAYER 2 State is MULTIPLE_FRAME_ESTABLISHED, and the Service Profile IDentifiers (SPIDs) are valid. If all of these conditions are satisfied, your problem is most likely not at Layers 1 or 2, and you should refer to the **debug isdn q931** command for further troubleshooting.

The **show isdn status** command displays the status of all ISDN interfaces or a specific ISDN interface. When troubleshooting ISDN BRI interfaces, it is necessary to first determine if the router can properly communicate with the telco ISDN switch. Once this has been verified, you can proceed on to higher level troubleshooting issues such as dialer interfaces, interesting traffic definitions, PPP negotiation, and authentication failures.

Note In certain parts of the world (notably in Europe) telco ISDN switches may deactivate Layer 1 or 2 when there are no active calls. Hence, when there are no active calls, **show isdn status** command will indicate that Layer 1 and 2 are down. However, when a call is placed Layers 1 and 2 will be activated. Make a test BRI call to verify whether the BRI is functioning. If the call succeeds, then no further ISDN troubleshooting is needed.

The configuration necessary for the router to communicate with the telco ISDN switch is fairly simple. You must have the ISDN switch type correctly configured for the BRI interface. Contact the telco to find out your circuit switch type.

You may also be required to have SPIDs configured. If you are connecting to a DMS-100 or NI-1 switch, you will most likely need to configure SPIDs. Most 5ESS switches do not require SPIDs. You should always contact your telco if you are unsure what switch type you are using.

Note If the telco informs you that SPIDs are not required, configure the interface as normal, skipping the **isdn spid1** and **isdn spid2** commands.

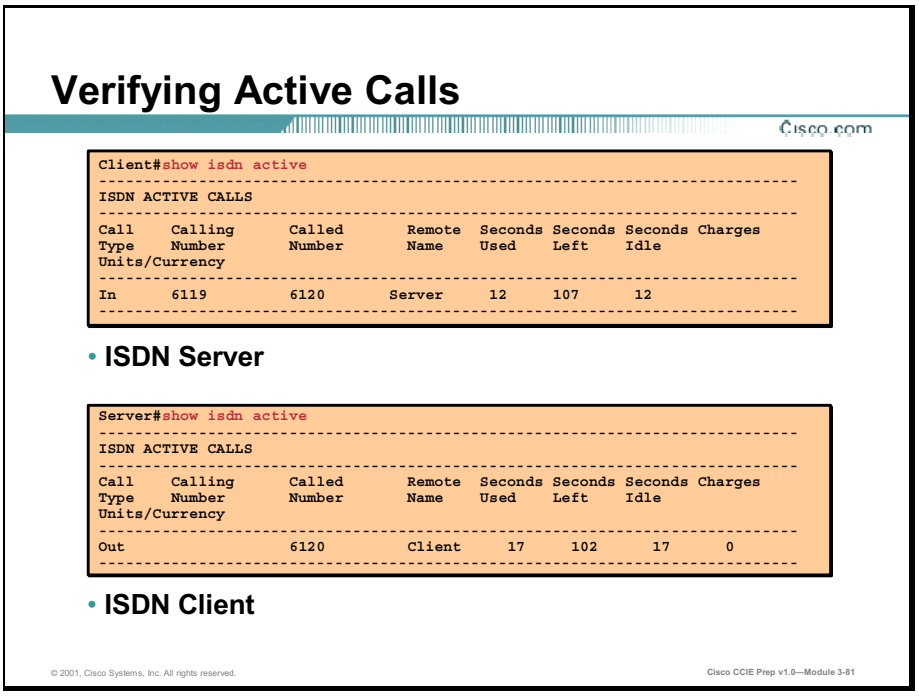

The **show isdn active** command displays information about the current call. This command can be used to verify that PPP callback was successfully completed. If callback is successful, **show isdn active** will show the call as incoming on the callback client and outgoing on the callback server.

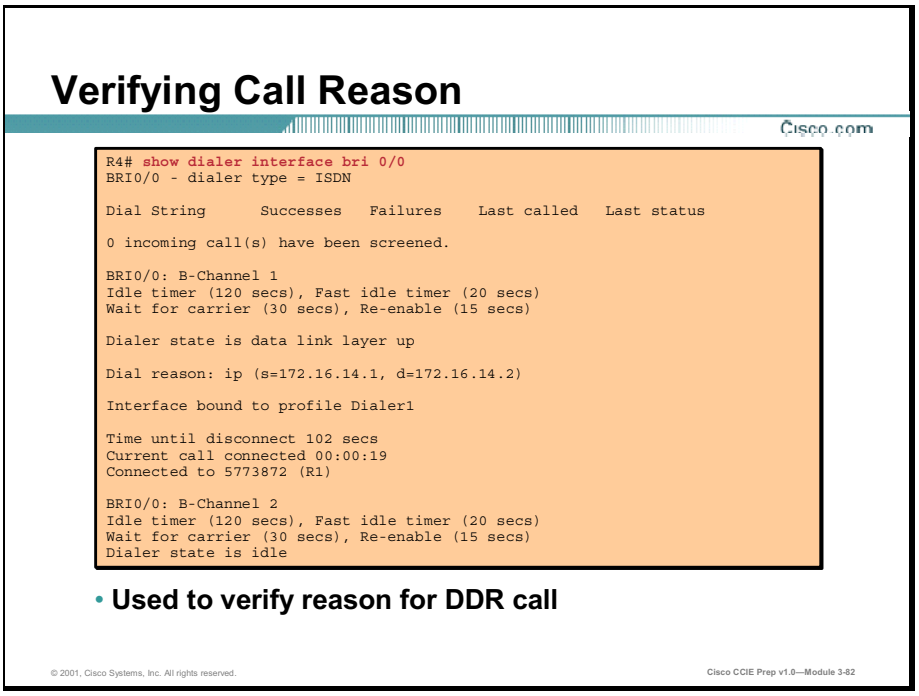

To display general diagnostic information for interfaces configured for Dial-on-Demand Routing (DDR), use the **show dialer** command in EXEC mode.

If you enter the **show dialer interface** command for the D channel of an ISDN BRI or PRI, the command output also displays the B channels. The **show dialer interface bri 0/0** command displays information of interfaces bri 0/0, bri 0/0:1, and bri 0/0:2. The **show dialer interface serial 0:23** command (for a channelized T1 line configured for ISDN PRI) displays information for serial interfaces 0:23, 0:0, 0:1, and so forth through 0:22.

If you have defined a dialer group that consists of the interfaces serial 0, serial 1, and bri 0/0, the **show dialer interface dialer 1** command displays information for interfaces bri 0/0, bri 0/0:1, bri 0/0:2, serial 1, and serial 0.

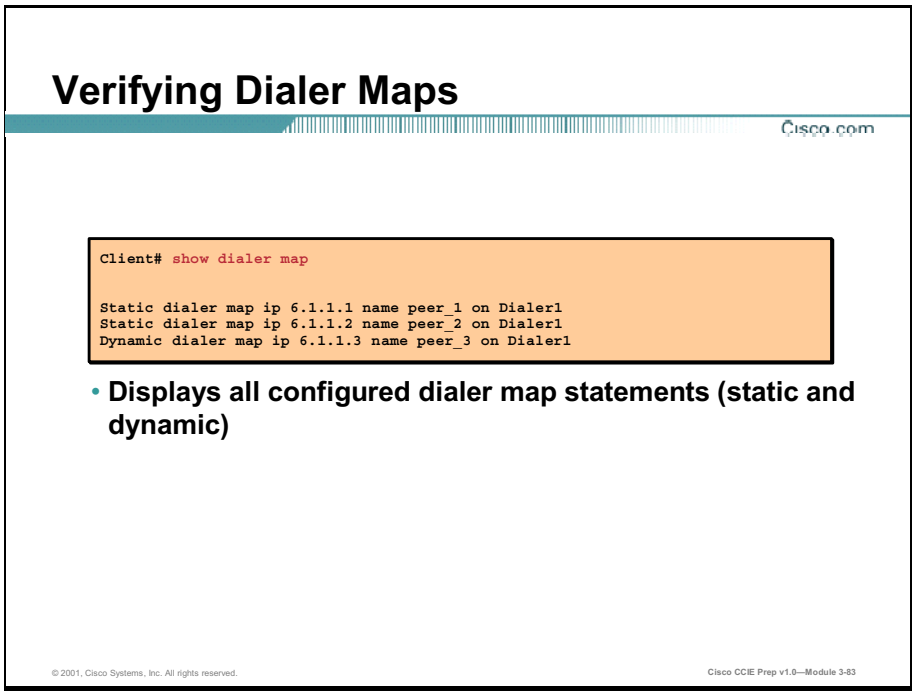

To display configured dialer maps, use the **show dialer map** command in EXEC mode.

The following table describes the significant fields in this output.

Table 3-2: Interpreting <show dialer map> Output

show dialer map fields	Description
`Static dialer map ip 6.1.1.1`	This is a statically configured dialer map to call the specified protocol address
`name peer_1`	Name of the remote peer
`on Dialer1`	The interface on which the static map is configured
`Dynamic dialer map ip 6.1.1.3`	Dialer map dynamically created when a peer is called

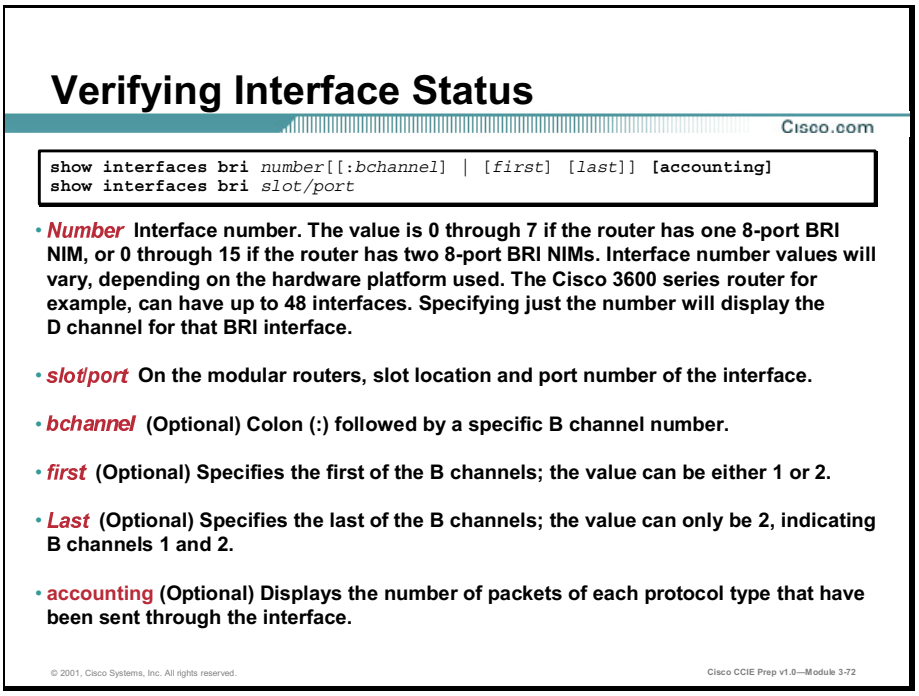

Use the **show interfaces bri** privileged EXEC command to display information about the BRI D channel or about one or more B channels.

Use either the ***:bchannel-number*** argument or the ***first*** or ***last*** arguments to display information about specified B channels.

Use the **show interfaces bri *number*** form of the command (without the optional ***:bchannel,*** or ***first*** and ***last*** arguments) to obtain D channel information.

Use the command syntax sample combinations in the following table to display the associated output.

Table 3-3: show interfaces bri Examples

Sample show interfaces bri Command Syntaxes	Displays
show interfaces	All interfaces in the router
show interfaces bri 2	Channel D for BRI interface 2
show interfaces bri 2:1	Channel B1 on BRI interface 2
show interfaces bri 2:2	Channel B2 on BRI interface 2
show interfaces bri 4 1	Channel B1 on BRI interface 4
show interfaces bri 4 2	Channel B2 on BRI interface 4
show interfaces bri 4 1 2	Channels B1 and B2 on BRI interface 4
show interfaces bri	Error message: "% Incomplete command."

Debug Commands

Debugging is sometimes necessary if **show** commands identify the proper configuration but an ISDN circuit continues to behave other than expected.

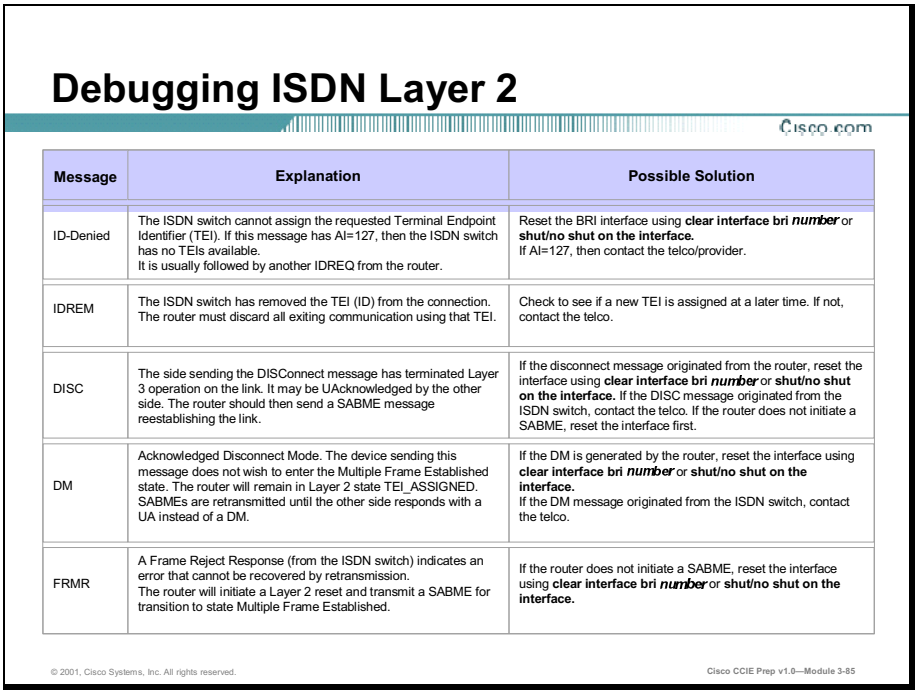

Use the **debug isdn q921** EXEC command to display data link layer (Layer 2) access procedures that are taking place at the router on the D channel (LAPD) of its Integrated Services Digital Network (ISDN) interface.

The ISDN data link layer interface provided by the router conforms to the user interface specification defined by ITU-T recommendation Q.921. The **debug isdn q921** command output is limited to commands and responses exchanged during peer-to-peer communication carried over the D channel. This debug information does not include data transmitted over the B channels that are also part of the router's ISDN interface. The peers (Data-Link layer entities and layer management entities on the routers) communicate with each other via an ISDN switch over the D channel.

Turn on **debug isdn q921** to see the messages that are transmitted from the router to the telco ISDN switch. You should then use the **clear interface bri *number*** to reset the BRI interface. This forces the router to renegotiate Layer 2 information with the telco ISDN switch.

Layer 2 problems cannot often be rectified at the customer site. However, Layer 2 debugs (or the interpretation of the debugs) can be provided to the telco for their reference. The **debug isdn q921** command output provides details on the Layer 2 transaction occurring between the ISDN switch and the router.

Pay attention to the direction of the messages. The debugs indicate whether the messages were generated by the router (indicated by TX ->) or if they were received by the router (indicated by

RX <--). In the example below, the first message (IDREQ) is sent by the router, while the second (IDASSN) is from the ISDN switch:

*Mar 1 00:03:46.976: ISDN BR0: **TX -> IDREQ** RI = 29609 AI = 127

*Mar 1 00:03:47.000: ISDN BR0: **RX <- IDASSN** RI = 29609 AI = 96

You can identify the source of the problem by following the direction of a particular message and the response. For example, if the telco ISDN switch unexpectedly sends a Layer 2 disconnect, the router will reset Layer 2 as well. This indicates that the problem lies with the telco ISDN switch.

Identifying Messages Indicating Layer 2 Problems

The router and the ISDN switch transmit and receive many Layer 2 messages. Most of the messages are normal and are used to verify normal operation. However, some messages can indicate Layer 2 problems. Though occasional resets may not affect service, if you observe extended periods of Layer 2 instability, you should take a closer look at the circuit.

The table shown on the previous page lists the **debug isdn q921** Layer 2 messages that indicate problems.

Here is an example of a received DISC message:

```
Jan 30 10:50:18.523: ISDN BR1/0: RX <- RRf sapi = 0 tei = 71 NR = 0
Jan 30 10:50:23.379: ISDN BR1/0: RX <- DISCp sapi = 0 tei = 71
Jan 30 10:50:23.379: %ISDN-6-Layer2DOWN: Layer 2 for Interface BR1/0,TEI 71
  changed to down
Jan 30 10:50:23.383: ISDN BR1/0: TX -> UAf sapi = 0 tei = 71
```

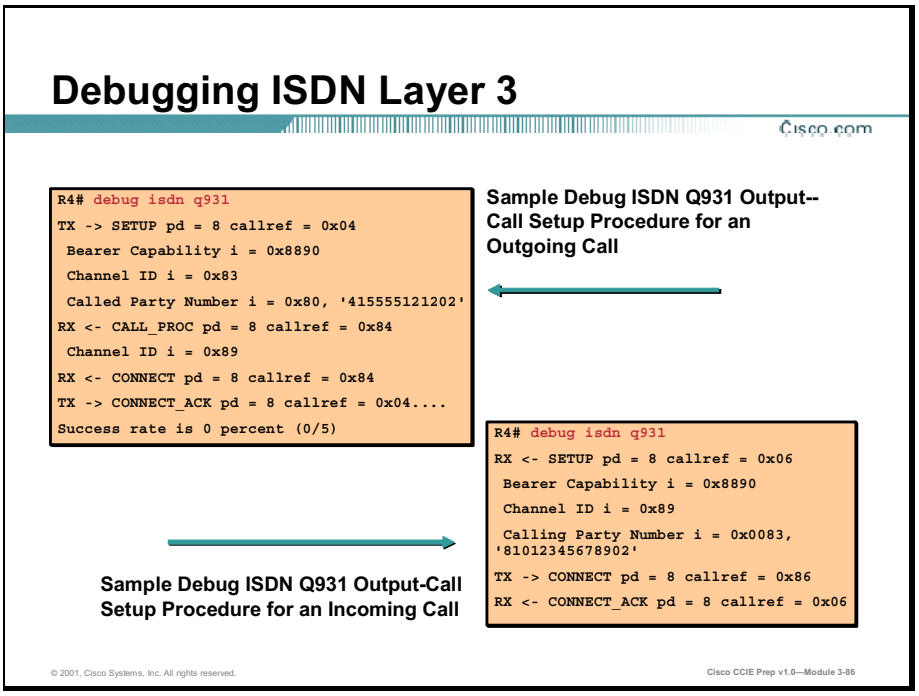

Use the **debug isdn q931** EXEC command to display information about call setup and teardown of ISDN network connections (Layer 3) between the local router (user side) and the network.

The ISDN network layer interface, provided by the router, conforms to the user interface specification defined by ITU-T recommendation Q.931. It is supplemented by other specifications such as switch type VN4. The router tracks only activities that occur on the user side, not the network side, of the network connection. The information displayed with the **debug isdn q931** command is limited to commands and responses exchanged during peer-to-peer communication carried over the D channel. This debug information does not include data transmitted over the B channels, which are also part of the router's ISDN interface. The peers (network layers) communicate with each other via an ISDN switch over the D channel.

A router can be the calling or called party of the ISDN Q.931 network connection call setup and teardown procedures. If the router is the calling party, the command displays information about an outgoing call. If the router is the called party, the command displays information about an incoming call.

You can use the **debug isdn q931** command with the **debug isdn events** and the **debug isdn q921** commands at the same time. The displays will be intermingled. Use the **service timestamps debug datetime msec** global configuration command to include the time with each message.

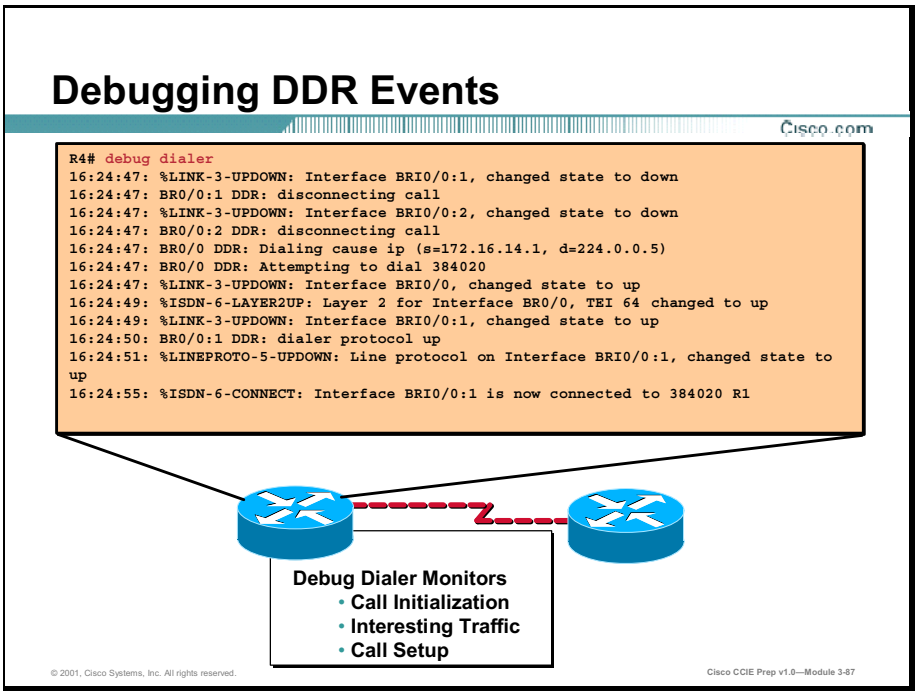

Use the **debug dialer** command to verify that the router is initiating a call properly and to verify that the DDR configuration is correct.

You can also use the **debug dialer** command to verify that the router is receiving interesting traffic and has the appropriate dialer map or dialer string to initiate the call.

Most messages are self-explanatory; however, messages that may need some explanation are described in the table below.

Table 3-4: Interpreting debug dialer Output

General debug dialer events Message Descriptions	Description
`Dialer0: Already xxx call(s) in progress on Dialer0, dialing not allowed`	This message occurs when the number of calls in progress (*xxx*) exceeds the maximum number of calls set on the interface.
`Dialer0: No free dialer - starting fast idle timer`	This message occurs when all the lines in the interface or rotary group are busy and a packet is waiting to be sent to the destination.
`BRI0: rotary group to xxx overloaded (yyy)`	This message occurs when the number dialer (*xxx*) exceeds the load set on the interface (*yyy*).
`BRI0: authenticated host xxx with no matching dialer profile`	This message occurs when no dialer profile matches *xxx*, the remote host's CHAP name or remote name.
`BRI0: Can't place call, verify configuration`	This message occurs when you have not set the dialer string or dialer pool on an interface.

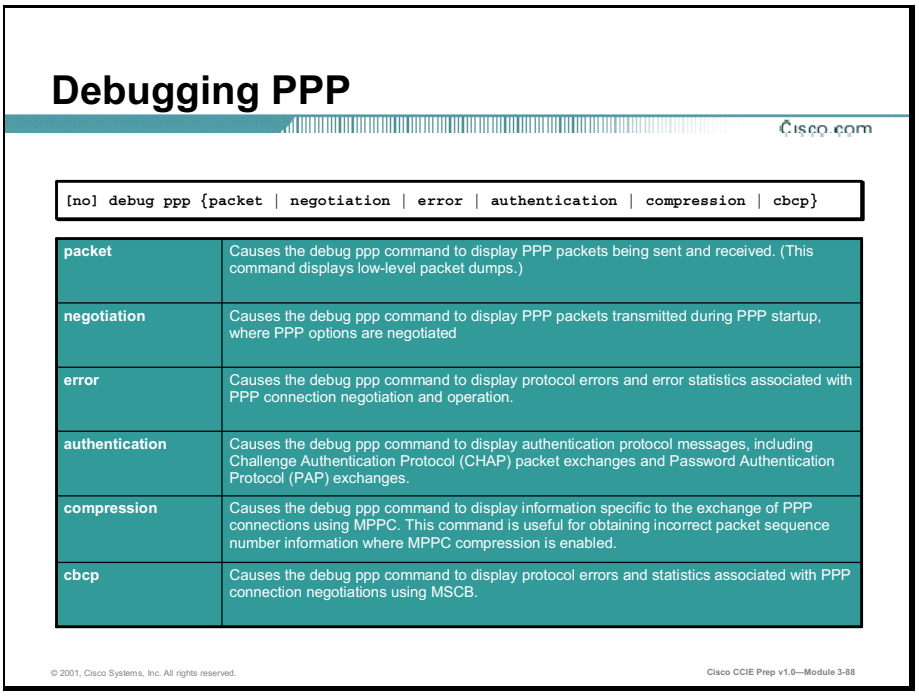

Use the **debug ppp** EXEC command to display information on traffic and exchanges in an internetwork implementing the Point-to-Point Protocol (PPP). The **no** form of this command disables debugging output.

Usage Guidelines

Use the **debug ppp** commands when trying to find the following:

- The Network Control Protocols (NCPs) that are supported on either end of a PPP connection
- Any loops that might exist in a PPP internetwork
- Nodes that are (or are not) properly negotiating PPP connections
- Errors that have occurred over the PPP connection
- Causes for CHAP session failures
- Causes for PAP session failures
- Information specific to the exchange of PPP connections using the Callback Control Protocol (CBCP), used by Microsoft clients
- Incorrect packet sequence number information where MPPC compression is enabled

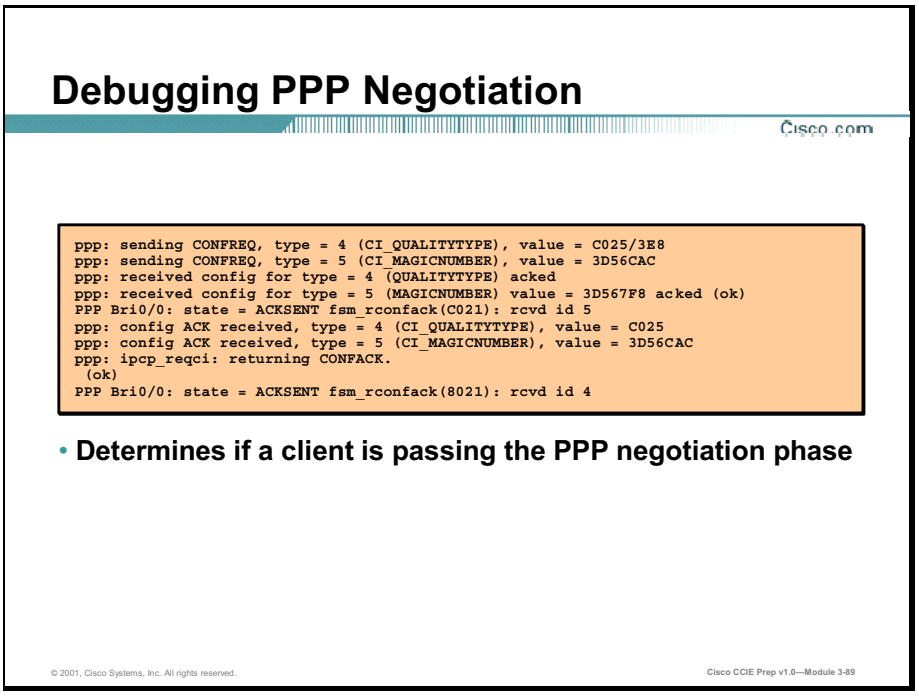

Use the **debug ppp negotiation** command to see if a client is passing PPP negotiation. This command is useful for verifying address negotiation.

The sample output shown is from the **debug ppp negotiation** command. This is a normal negotiation, where both sides agree on Network Control Program (NCP) parameters. In this case, protocol type IP is proposed and acknowledged.

The following table describes significant fields in the output.

Table 3-5: Interpreting debug ppp negotiation Output

debug ppp negotiation Field Descriptions	Description
`ppp`	This is PPP debugging output.
`sending CONFREQ`	The router sent a configuration request.
`type = 4 (CI_QUALITYTYPE)`	The type of LCP configuration option that is being negotiated and a descriptor. A type value of 4 indicates Quality Protocol negotiation; a type value of 5 indicates Magic Number negotiation.
`value = C025/3E8`	For Quality Protocol negotiation, indicates NCP type and reporting period. In the example, C025 indicates LQM; 3E8 is a hexadecimal value translating to about 10 seconds (in hundredths of a second).
`value = 3D56CAC`	For Magic Number negotiation, indicates the Magic Number being negotiated.
`received config`	The receiving node has received the proposed option negotiation for the indicated option type.
`acked`	Acknowledgment and acceptance of options.
`state = ACKSENT`	Specific PPP state in the negotiation process.

debug ppp negotiation Field Descriptions	Description
ipcp_reqci	IPCP notification message; sending CONFACK.
fsm_rconfack (8021)	The procedure fsm_rconfack processes received CONFACKs, and the protocol (8021) is IP.

The following two lines of syntax indicate that the router is trying to bring up LCP and will use the indicated negotiation options (Quality Protocol and Magic Number). The value fields are the values of the options themselves. C025/3E8 translates to Quality Protocol LQM. 3E8 is the reporting period (in hundredths of a second). 3D56CAC is the value of the Magic Number for the router.

```
ppp: sending CONFREQ, type = 4 (CI_QUALITYTYPE), value = C025/3E8
ppp: sending CONFREQ, type = 5 (CI_MAGICNUMBER), value = 3D56CAC
```

The next two lines indicate that the other side negotiated for options 4 and 5 as requested and acknowledged both. If the responding end does not support the options, a CONFREJ is sent by the responding node. If the responding end does not accept the value of the option, a CONFNAK is sent with the value field modified.

```
ppp: received config for type = 4 (QUALITYTYPE) acked
ppp: received config for type = 5 (MAGICNUMBER) value = 3D567F8 acked (ok)
```

The next three lines indicate that the router received a CONFACK from the responding side and displays accepted option values. Use the rcvd id field to verify that the CONFREQ and CONFACK have the same id field.

```
PPP Bri0/0: state = ACKSENT fsm_rconfack(C021): rcvd id 5
ppp: config ACK received, type = 4 (CI_QUALITYTYPE), value = C025
ppp: config ACK received, type = 5 (CI_MAGICNUMBER), value = 3D56CAC
```

The next line indicates that the router has IP routing enabled on this interface and that the IPCP NCP negotiated successfully:

```
ppp: ipcp_reqci: returning CONFACK.
```

In the last line, the router's state is listed as ACKSENT.

```
PPP Bri0/0: state = ACKSENT fsm_rconfack(C021): rcvd id 5\
```


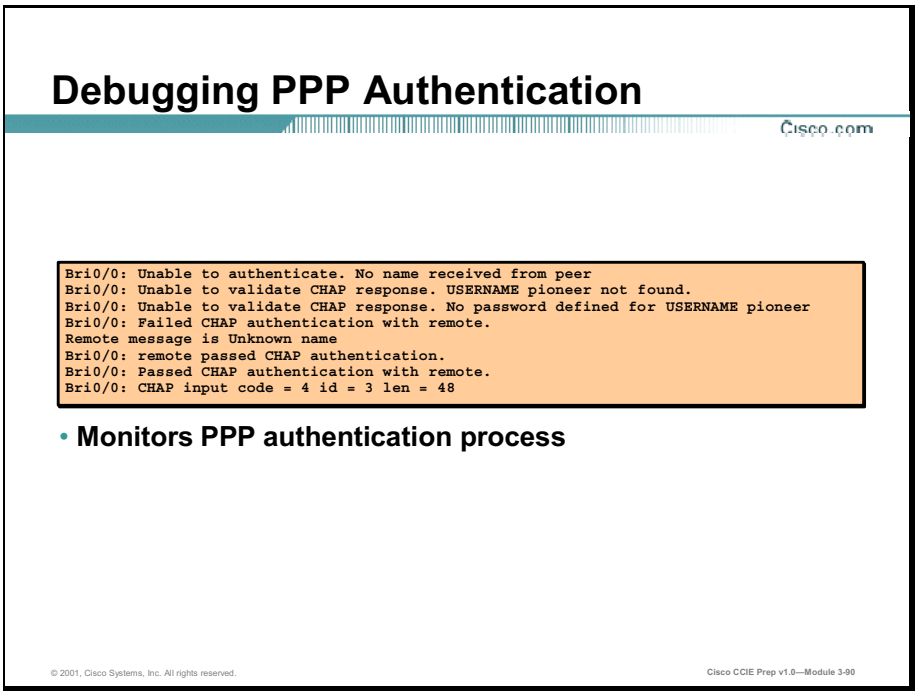

Shown here is sample output from the **debug ppp authentication** command. Use this command to determine why authentication is failing between two peer routers.

In general, these messages are self-explanatory. Fields that can show optional output are outlined in the following table.

Table 3-6: Interpreting debug ppp authentication Output

debug ppp authentication Field Descriptions Field	Description
`Bri0/0`	Interface number associated with this debugging information and CHAP access session in question
`USERNAME pioneer not found.`	The name pioneer in this example is the name received in the CHAP response. The router looks up this name in the list of usernames that are configured for the router
`Remote message is Unknown name`	The following messages can appear: No name received to authenticate Unknown name No secret for given name Short MD5 response received MD compare failed

debug ppp authentication Field Descriptions Field	Description
`code = 4`	Specific CHAP type packet detected. Possible values are as follows: 1 = Challenge 2 = Response 3 = Success 4 = Failure

Summary

This section summarizes the key points discussed in this lesson.

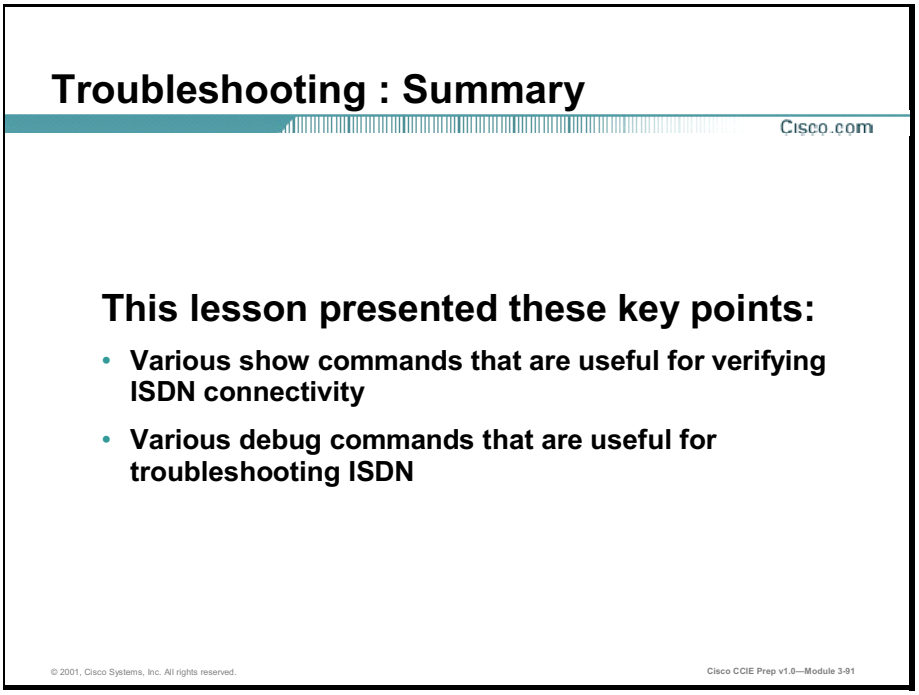

Next Steps

After completing this lesson, go to:

- Ethernet Switching Technologies

References

For additional information, refer to these resources:

- http://www.cisco.com/warp/public/779/smbiz/service/troubleshooting/ts_isdn.htm

Lesson Assessment (Quiz)

Q1) Which **show** command is useful to view the ISDN information for layers 1, 2 and 3?

A) **show isdn q921**

B) **show isdn q931**

C) **show isdn status**

D) **show isdn active**

Q2) Which **show** command can be used to show detailed information about calls in progress?

A) **show isdn active**

B) **show isdn q931**

C) **show isdn status**

D) None of the above

Q3) What **show** command is useful for the verification of DDR setup?

A) **show isdn q931**

B) **show isdn active**

C) **show interfaces bri**

D) **show dialer interface**

Q4) Which ISDN **debug** command displays data link layer information?

A) **debug isdn q921**

B) **debug isdn q931**

C) **debug dialer**

D) None of the above

Q5) Which ISDN **debug** command is most appropriate for verifying DDR operation?

A) **debug isdn q921**

B) **debug dialer**

C) **debug isdn q931**

D) None of the above

4

Catalyst 3550 Switching

Overview

This module will focus on the configuration of the two 3550 Catalyst switches in your CCIE equipment rack.

Upon completing this module, you will be able to:

- Configure basic features on the Catalyst 3550, such as VTP and VLANs.
- Configure the different types of interfaces available on the Catalyst 3550, such as Access Ports, Trunk Ports, Tunnel Ports, Router Ports, and Switched Virtual Interfaces (SVI).
- Fine tune Spanning Tree operation on the Catalyst 3550
- Monitor and analyze network traffic using the Switched Port Analyzer (SPAN) and Remote Switched Port Analyzer (RSPAN) features.
- Configure Fallback Bridging to bridge non-IP traffic.

Outline

The module contains these lessons:

- Catalyst 3550 Basic Configuration
- Catalyst 3550 Interface Configuration
- Catalyst 3550 Advanced Configuration

Catalyst 3550 Basic Configuration

Overview

The Catalyst 3550 is an IOS-based Multilayer switch. The Catalyst 3550 can operate as either a Layer 2 or Layer 3 device or it can act as a Multilayer switch incorporating both Layer 2 and Layer 3 features. In order for the Catalyst 3550 to support Layer 3 features, such as routing, it must run the Enhanced Multilayer Image (EMI). The CCIE candidate's equipment rack will include two Catalyst 3550 switches running the EMI image.

Importance

The Catalyst 3550 Switch is the backbone of the LAN, the device through which all LAN routers communicate with each other. Correct configuration of the basic features (VTP and VLANs) on Catalyst 3550 switch is critical in the CCIE Lab.

Objectives

Upon completing this lesson, you will be able to configure:

- Assign an IP address to the switch's management interface and allow remote management of the Catalyst 3550
- Configure the VLAN Trunking Protocol (VTP) to allow the creation of VLANs on the Catalyst 3550
- Configure Virtual LANs (VLANs) to segment users and traffic
- Verify the correct configuration of VTP and VLANs using the available **show** commands

Learner Skills and Knowledge

To fully benefit from this lesson, you must have these prerequisite skills and knowledge:

- Passed the CCIE written qualification exam
- Completed the Building Cisco Multilayer Switched Networks (BCMSN) course, or have the equivalent knowledge

Outline

This lesson includes these sections:

- Overview
- Management Interface Configuration
- VTP Configuration
- VLAN Configuration
- Troubleshooting VTP and VLANs
- Summary
- Lesson Assessment (Quiz)

Management Interface Configuration

This section describes the initial configuration of the Catalyst 3550 switch. For example, assigning the switch an IP address and default gateway.

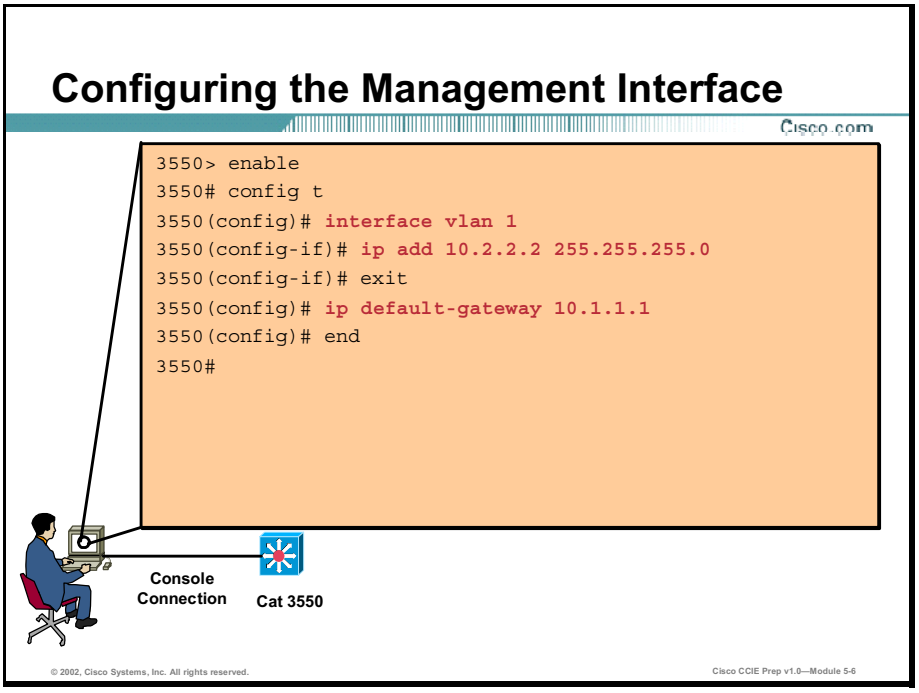

Before you can configure the switch, you need connectivity to it via the console port. Make sure the terminal connected to the console port is configured as follows: 9600 baud, 8 data bits, no parity, and 1 stop bit.

When the switch boots, its management interface's address is set to 0.0.0.0 (the default on a new switch or after the configuration is cleared). Upon bootup, the switch attempts to obtain an IP address using Dynamic Host Configuration Protocol (DHCP). If required, a Cisco router running the DHCP service could support this.

On the Catalyst 3550, VLAN1 is reserved for the management interface. The following table outlines the steps to manually assign an IP address to the management interface of the Catalyst 3550.

Table 4-1: Manually Assign an IP Address

Command	Purpose
3550(config)# **interface vlan 1**	Enter interface configuration mode for VLAN1.
3550(config-if)#**ip address** *ip-address subnet-mask*	Enter the IP address and subnet mask.
3550(config)# **ip default-gateway** *ip-address*	Enter the IP address of the next-hop router that should be used when traffic is destined for a host that is not on the same VLAN. Once the default gateway is configured, the switch has connectivity to the remote networks with which a host needs to communicate. **Note** If your switch is configured to route IP, it does not need to have a default gateway set.

VTP Configuration

This chapter describes how to use the VLAN Trunking Protocol (VTP) and the VLAN database for creating and managing VLANs.

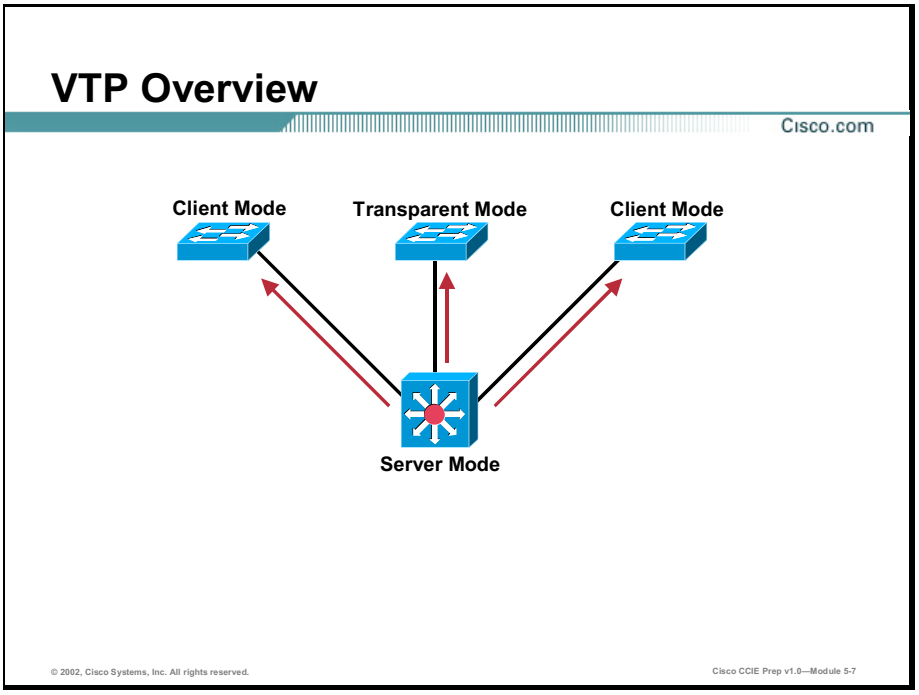

VTP is a Layer 2 messaging protocol that maintains VLAN configuration consistency by managing the addition, deletion, and renaming of VLANs on a network-wide basis. VTP minimizes misconfigurations and configuration inconsistencies that can cause several problems, such as duplicate VLAN names, incorrect VLAN-type specifications, and security violations.

Before you create VLANs, you must decide whether to use VTP in your network. Using VTP, you can make configuration changes centrally on one or more switches and have those changes automatically communicated to all the other switches in the network. Without VTP, you cannot send information about VLANs to other switches. VTP configuration information is saved in the VTP VLAN database. Catalyst switches can support VTP in one of three modes: Server, Client, and Transparent.

- Server: Allows you to create, modify, and delete VLANs and specify other configuration parameters (such as VTP version and VTP pruning) for the entire VTP domain. VTP servers advertise their VLAN configuration to other switches in the same VTP domain and synchronize their VLAN configuration with other switches based on advertisements received over trunk links. VTP server is the default mode.

- Client: Behaves the same way as VTP servers, except that you cannot create, change, or delete VLANs on a VTP client.

- Transparent: Switches do not participate in VTP. A VTP transparent switch does not advertise its VLAN configuration and does not synchronize its VLAN configuration based on received advertisements. However, in VTP version 2, transparent switches do forward VTP advertisements that they receive out their trunk ports. Transparent mode is required if you want to configure a switch to support extended range VLANs.

On the Catalyst 3550 you can configure VTP in one of two ways: using the **vtp** global configuration command or using the **vtp** commands available in VLAN configuration mode.

Note	VTP only learns about normal-range VLANs (VLAN IDs 1 to 1005). Extended-range VLANs (VLAN IDs greater than 1005) are not supported by VTP or stored in the VTP VLAN database.

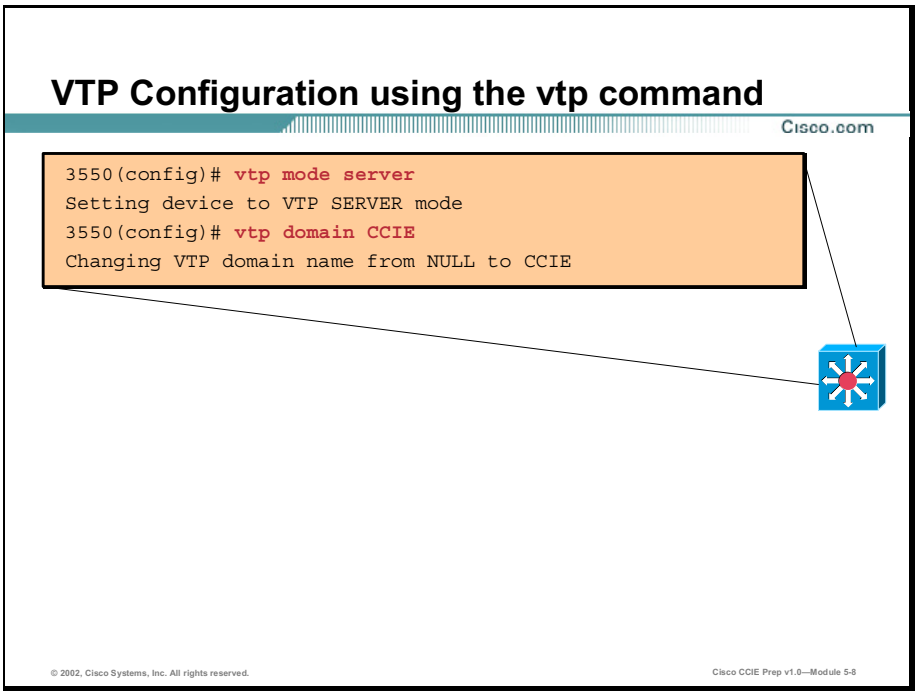

You can use the **vtp** global configuration command to set the VTP domain, mode, password, version, VTP file name, and to disable or enable VTP pruning. The information entered with the **vtp** global configuration command is saved in the VTP VLAN database. When a switch is in VTP server mode, you can change the VLAN configuration and have it propagated throughout the network. In the CCIE lab, you will mostly likely configure at least one of your switches as a VTP Server.

Use the steps outlined in the following table to configure the Catalyst 3550 switch as a VTP server:

Table 4-2: Configure switch as a VTP Server

Command	Purpose
3550(config)**# vtp mode server**	Configure the switch for VTP server mode (default).
3550(config)# **vtp domain** *domain-name*	Configure the VTP administrative-domain name. The name can be from 1 to 32 characters. All switches operating in VTP server or client mode under the same administrative responsibility must be configured with the same domain name.

Note When you configure a VTP domain name, it cannot be removed; you can only reassign a switch to a different VTP domain.

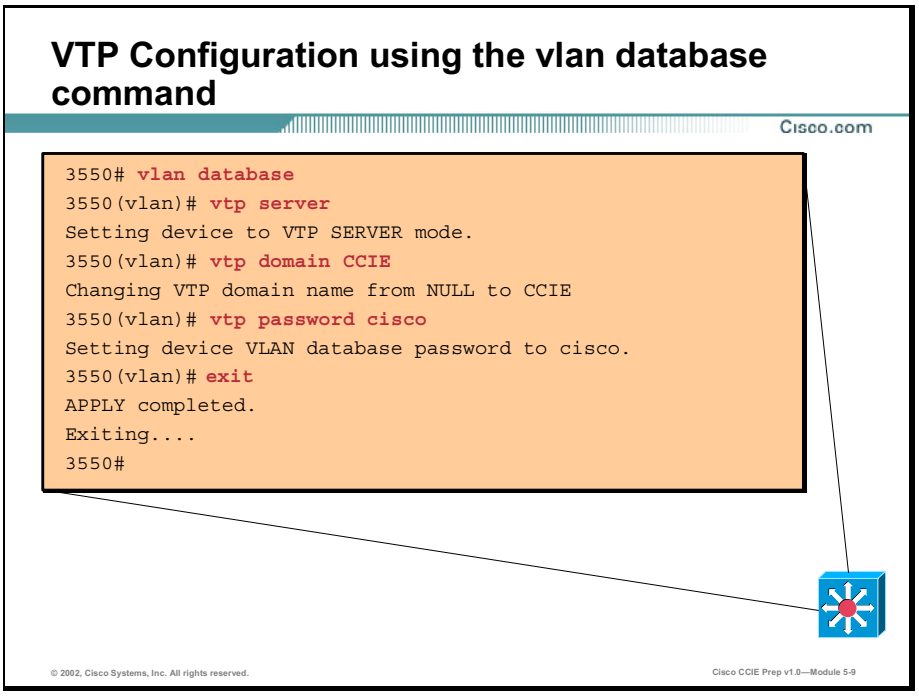

You can also configure all VTP parameters in VLAN configuration mode, which you access by entering the **vlan database** command. When you enter the **exit** command in VLAN configuration mode, it applies all the commands that you entered and updates the VLAN database.

Note	Configuring VTP via the **vlan database** command is the preferred method as some of the advanced settings, such as setting a VTP password, enabling VTP version 2, and enabling VTP pruning, can only be done in the vlan database (vlan) configuration mode.

Use the steps outlined in the following table to use VLAN configuration mode and configure the switch as a VTP server:

Table 4-3: VLAN configuration mode

Command	Purpose
3550**# vlan database**	Enter VLAN configuration mode.
3550(vlan)# **vtp server**	Configure a VTP administrative-domain name. The name can be from 1 to 32 characters. All switches operating in VTP server or client mode under the same administrative responsibility must be configured with the same domain name.
3550(vlan)**# vtp password** *password*	(Optional) Set a password for the VTP domain. The password can be from 8 to 64 characters. If you configure a VTP password, the VTP domain does not function properly if you do not assign the same password to each switch in the domain.
3550(vlan)# **exit**	Update the VLAN database, propagate changes throughout the administrative domain, and return to privileged EXEC mode.

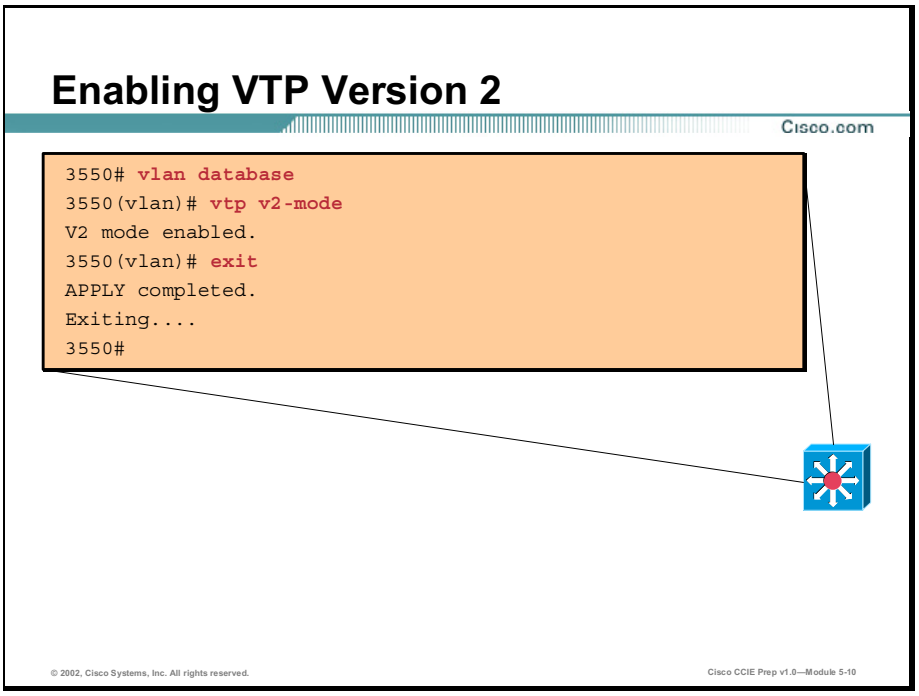

VTP version 2 is disabled by default on VTP version 2-capable switches. When you enable VTP version 2 on a switch, every VTP version 2-capable switch in the VTP domain enables version 2. You can only configure the version on switches in VTP server or transparent mode.

Note	VTP version 1 and VTP version 2 are not interoperable on switches in the same VTP domain. Every switch in the VTP domain must use the same VTP version. Do not enable VTP version 2 unless every switch in the VTP domain supports version 2.

Use the steps outlined in the following table to enable VTP version 2 on the Catalyst 3550:

Table 4-4: VLAN version 2

Command	Purpose
3550(vlan)# **vtp v2-mode**	Enable VTP version 2 on the switch. VTP version 2 is disabled by default on VTP version 2-capable switches.

To disable VTP version 2, use the **no vtp v2-mode** vlan database (vlan) configuration command.

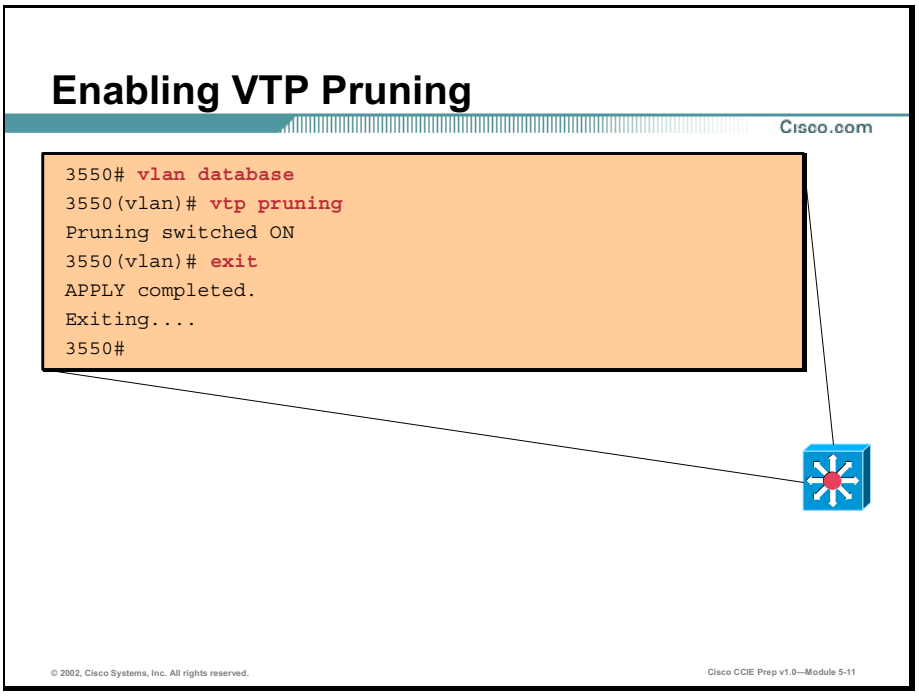

Pruning increases available bandwidth by restricting flooded traffic to those trunk links that the traffic must use to access the destination devices. You can only enable VTP pruning on a switch in VTP server mode. Pruning is supported with VTP version 1 and version 2. If you enable pruning on the VTP server, it is enabled for the entire VTP domain.

Only VLANs included in the pruning-eligible list can be pruned. By default, VLANs 2 through 1001 are pruning eligible on trunk ports. Reserved VLANs and extended-range VLANs cannot be pruned.

Use the steps outlined in the following table to enable VTP pruning in the VTP domain:

Table 4-5: VLAN Pruning

Command	Purpose
3550(vlan)# **vtp pruning**	Enable pruning in the VTP administrative domain. By default, pruning is disabled. You need to enable pruning on only one switch in VTP server mode.

To disable VTP pruning, use the **no vtp pruning** vlan database (vlan) configuration command.

VLAN Configuration

Once the switch is properly configured for VTP, you can create, modify, and delete VLANs on the switch (unless you configured the switch as a VTP client). The default Ethernet VLAN is VLAN 1. By default, all switch ports are assigned to VLAN 1. Once VTP is configured, you can create additional VLANs and assign specific switch ports to those VLANs.

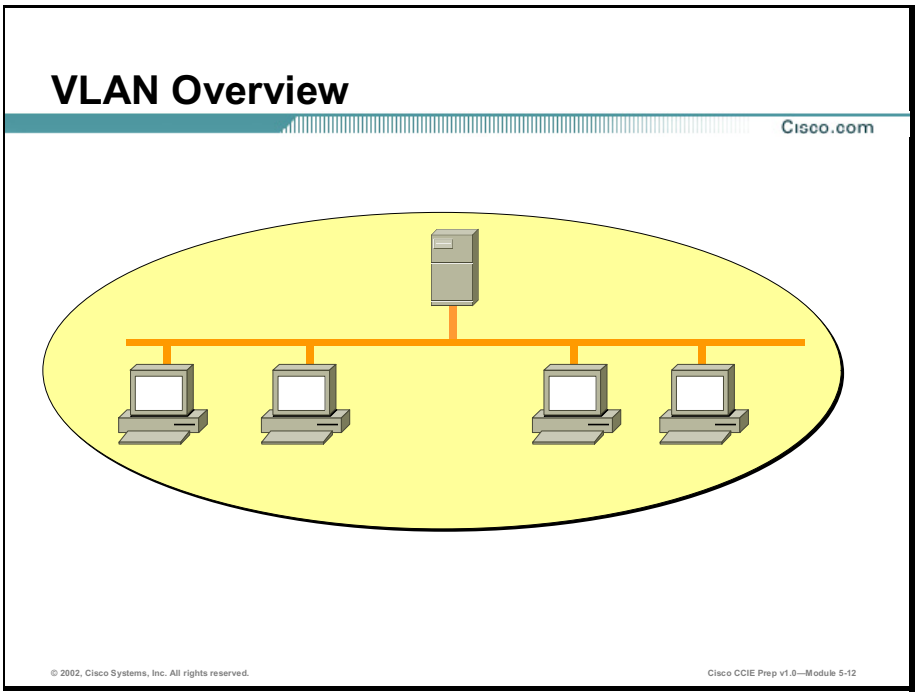

A VLAN is a switched network that is logically segmented by function, project team, or application, without regard to the physical locations of the users. VLANs have the same attributes as physical LANs, except that you can group end stations together even if they are not physically located on the same LAN segment. Any switch port can belong to a VLAN, and unicast, broadcast, and multicast packets are forwarded and flooded only to end stations in the VLAN. Each VLAN is considered a logical network, and packets destined for stations that do not belong to the same VLAN must be forwarded through a router or other Layer 3 engine.

You can configure normal-range VLANs (with VLAN IDs 1 to 1005) by using one of two configuration modes:

- VLAN Configuration in (config-vlan) Mode

You can access (config-vlan) mode by entering the **vlan** *vlan-id* global configuration command.

- VLAN Configuration in VLAN Configuration Mode

You can access VLAN configuration mode by entering the **vlan database** privileged EXEC command.

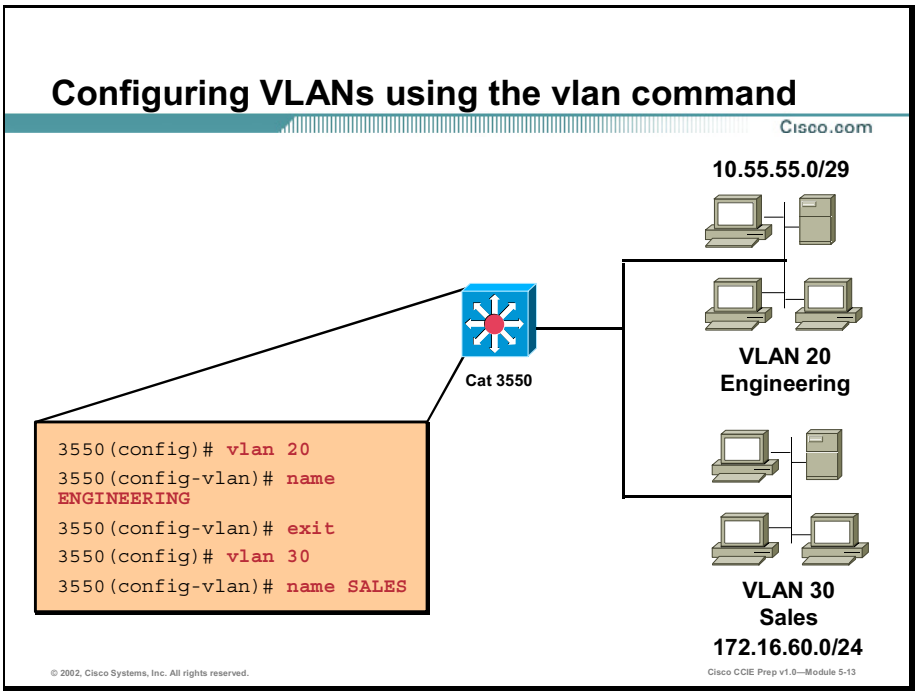

Each Ethernet VLAN in the VLAN database has a unique, 4-digit ID that can be a number from 1 to 1001. VLAN IDs 1002 to 1005 are reserved for Token Ring and FDDI VLANs. To create a normal-range VLAN to be added to the VLAN database, assign a number and name to the VLAN.

Note	When the switch is in VTP transparent mode, you can assign VLAN IDs greater than 1006 (Extended Range VLANs), but they are not added to the VLAN database.

Use the steps outlined in the following table to use (config-vlan) mode to create or modify an Ethernet VLAN:

Table 4-6: Ethernet VLAN

Command	Purpose
3550(config)# **vlan** *vlan-id*	Enter a new VLAN ID to create a VLAN, or enter an existing VLAN ID to modify a VLAN.
3550(config-vlan)# **name** *vlan-name*	(Optional) Enter a name for the VLAN. If no name is entered for the VLAN, the default is to append the *vlan-id* with leading zeros to the word VLAN. For example, VLAN0004 is a default VLAN name for VLAN 4.

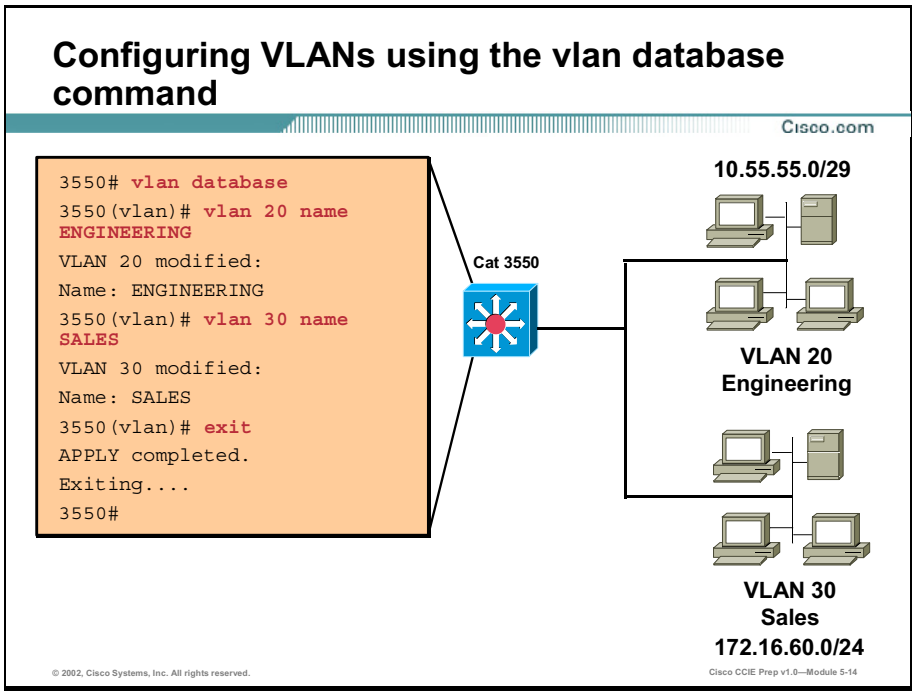

Use the steps outlined in the following table to use VLAN configuration mode to create or modify an Ethernet VLAN:

Table 4-7: Ethernet VLAN

Command	Purpose
3550# **vlan database**	Enter VLAN configuration mode.
3550(vlan)# **vlan** *vlan-id* **name** *vlan-name*	Add an Ethernet VLAN by assigning a number to it. The range is 1 to 1001; do not enter leading zeros. If no name is entered for the VLAN, the default is to append the *vlan-id* with leading zeros to the word VLAN. For example, VLAN0004 is a default VLAN name for VLAN 4.
3550(vlan)# **exit**	Updates the VLAN database, propagate changes throughout the VTP administrative domain, and returns to privileged EXEC mode.

Troubleshooting VTP and VLANs

This section covers the commands that can be used to troubleshoot VTP and VLAN problems.

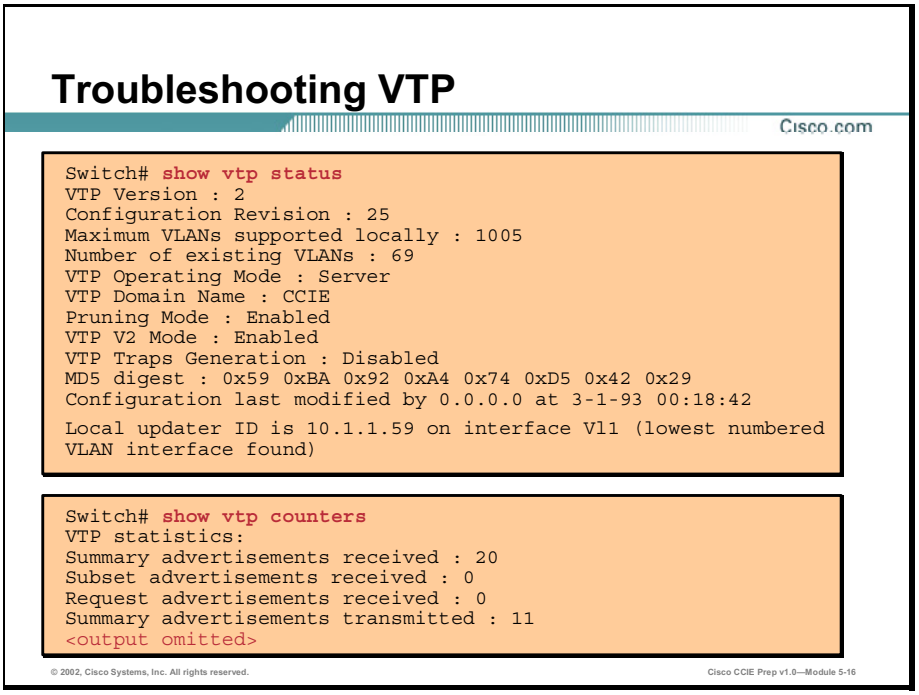

VTP problems usually arise because of inconsistencies in the VTP database on the different switches throughout your network. You can verify VTP synchronization by displaying VTP configuration information: the domain name, the current VTP revision, and the number of VLANs on the different switches throughout your network. You can also display statistics about the number of advertisements sent and received by the switch.

The following table shows the commands used for troubleshooting VTP:

Table 4-8: Troubleshooting VTP

Command	Purpose
3550# **show vtp status**	Displays the VTP switch configuration information.
3550# **show vtp counters**	Displays counters about VTP messages that have been sent and received.

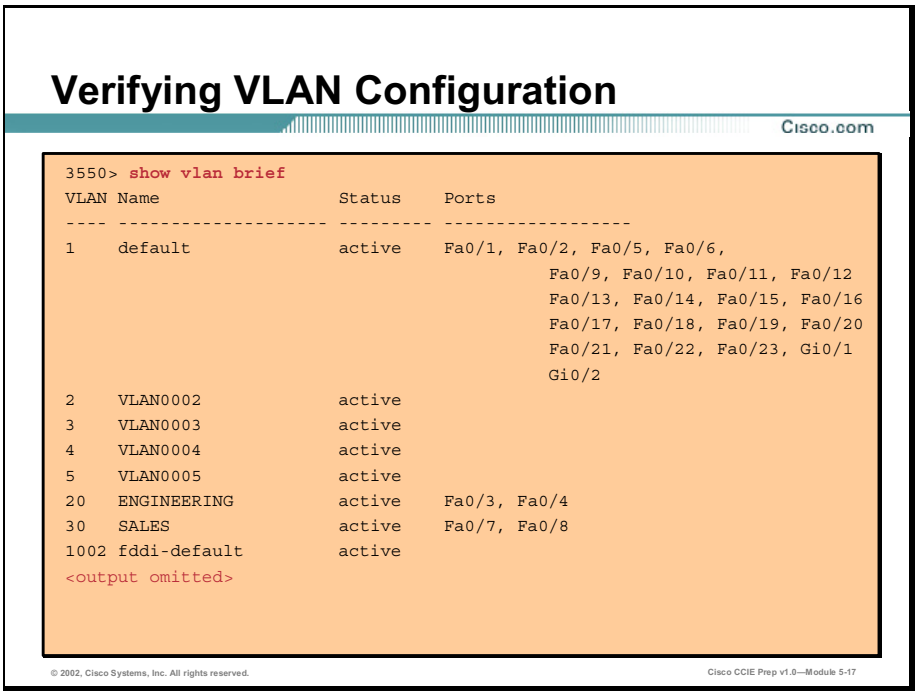

In much the same manner as VTP, VLAN problems usually arise when there is inconsistent VLAN information on the different switches in the network. You can use the **show vlan** command to display a list of all VLANs on each switch, including extended-range VLANs. The display includes VLAN status, ports, and configuration information. To view normal-range VLANs in the VLAN database (1 to 1005,) use the **show** VLAN configuration command (accessed by entering the **vlan database** privileged EXEC command).

The following table lists the commands used for verifying VLAN configuration and status:

Table 4-9: VLAN Configuration

VLAN Monitoring Commands	Purpose
3550(vlan)# show	Displays the status of all VLANs in the VLAN database.
3550(vlan)# show current [*vlan-id*]	Displays the status of all or the specified VLAN in the VLAN database.
3550# show running-config vlan	Displays the running configuration all or a range of VLANs on the switch.
3550# show vlan [id *vlan-id*]	Displasy parameters for all VLANs or the specified VLAN on the switch.
3550# show vlan brief	Displays a brief summary of all VLANs configured on the switch.

Summary

This section summarizes the key points discussed in this lesson.

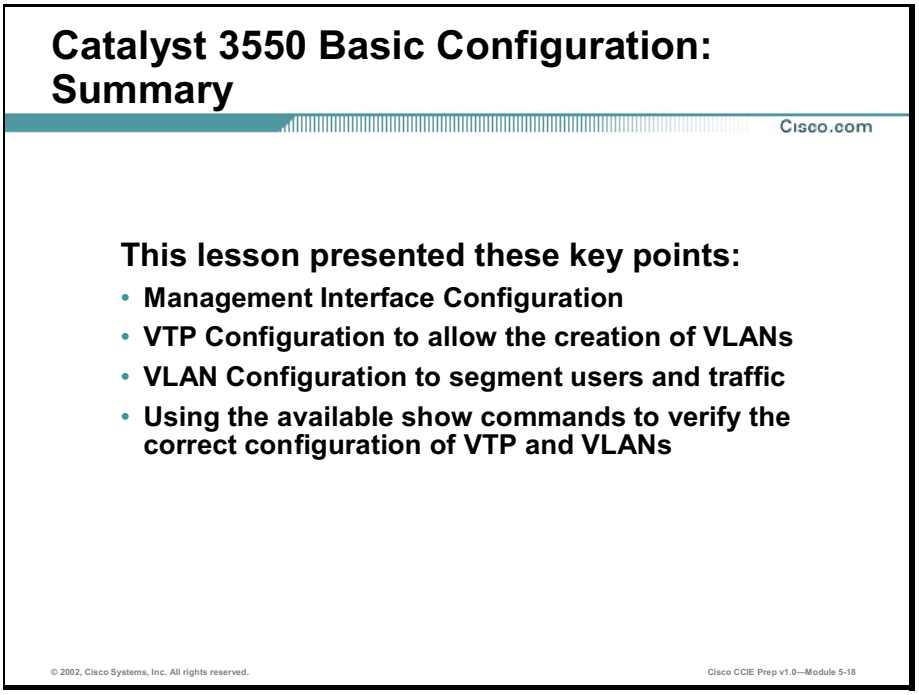

Next Steps

After completing this lesson, go to:

- Catalyst 3550 Interface Configuration

References

For additional information, refer to these resources:

- *CCIE Professional Development : Cisco Lan Switching* by Kennedy Clark

Lesson Assessment (Quiz)

Q1) The management interface on the Catalyst 3550 belongs to which VLAN by default?

A) VLAN 1

B) All VLANs (it is a trunk port)

C) None – you must create a SVI for VLAN 1 first

D) VLAN 1005

Q2) Which VTP mode should you use if you wish to configure Extended Range VLANs?

A) Server

B) Client

C) Transparent

D) The Catalyst 3550 does not support Extended Range VLANs

Q3) When creating VLANs using the vlan database command, when are your changes actually made to the VLAN database and propagated to other switches in the VTP domain?

A) As soon as the VLAN is created

B) Once you give the VLAN a name

C) Once the switch is rebooted

D) When you enter the **exit** command to go back to privileged exec mode

Q4) What command can be used to obtain a brief summary of all of the VLANs configured on the switch?

show vlan brief

Catalyst 3550 Interface Configuration

Overview

The Catalyst 3550 running the Enhanced Multilayer Image (EMI) supports many different types of interfaces. Some of these interfaces are actual physical interfaces, such as switch ports and router ports, while others are logical interfaces, such as Switched Virtual Interfaces (SVI) and EtherChannel Port Groups. This section will discuss the differences between the different types of interfaces and their respective configurations.

Importance

In order to successfully configure many of the features on the Catalyst 3550, including VLANs, 802.1Q and ISL Trunking, 802.1Q and Layer 2 Tunneling, EtherChannel, Fallback Bridging, you understand the different types of interfaces the Catalyst 3550 supports and how to configure them.

Objectives

Upon completing this lesson, you will be able to configure:

- List the three different types of switch ports that the Catalyst 3550 supports
- Statically configure Access Ports for VLANs
- Configure Trunk Ports for 802.1Q and ISL trunking
- Configure Tunnel Ports for 802.1Q and Layer 2 Protocol Tunneling
- Configure Layer 3 Interfaces (Router Ports and SVIs)
- Configure EtherChannel for Layer 2 and Layer 3 Interfaces

Learner Skills and Knowledge

To fully benefit from this lesson, you must have these prerequisite skills and knowledge:

- Passed the CCIE written qualification exam
- Completed the Building Cisco Multilayer Switched Networks (BCMSN) course, or have the equivalent knowledge

Outline

This lesson includes these sections:

- Overview
- Overview of Switch Ports
- Access Port Configuration
- Trunk Port Configuration
- Tunnel Port Configuration
- Layer 3 Interfaces
- General Interface Commands
- EtherChannel
- Summary
- Lesson Assessment (Quiz)

Overview of Switch Ports

This section will provide an overview of the different types of switch ports available on the Catalyst 3550.

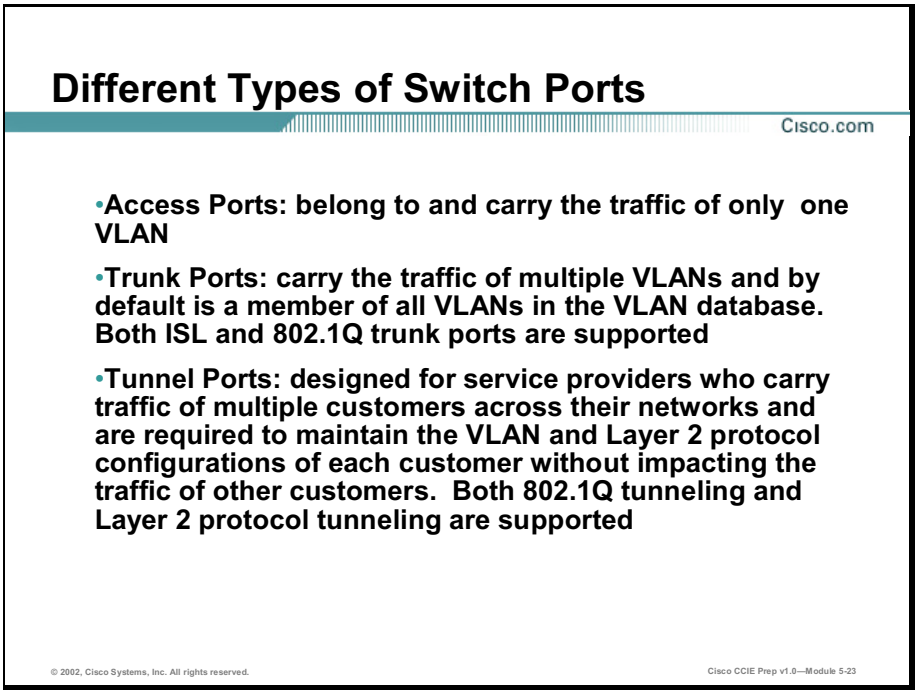

Switch ports are Layer 2-only interfaces associated with a physical port. A switch port can be an access port, a trunk port, or a tunnel port. You can manually configure a port as an access port or trunk port or let the Dynamic Trunking Protocol (DTP) operate on a per-port basis to determine if a switch port should be an access port or a trunk port by negotiating with the port on the other end of the link. You must manually configure tunnel ports as part of an asymmetric link connected to an 802.1Q trunk port. Switch ports are used for managing the physical interface and associated Layer 2 protocols and do not handle routing or bridging.

Access Ports

An access port belongs to and carries the traffic of only one VLAN. Traffic is received and sent in native formats with no VLAN tagging. Traffic arriving on an access port is assumed to belong to the VLAN assigned to the port.

Trunk Ports

A trunk port carries the traffic of multiple VLANs and by default is a member of all VLANs in the VLAN database. Two types of trunk ports are supported:

- ISL trunk port - All received packets are expected to be encapsulated with an ISL header, and all transmitted packets are sent with an ISL header. Native (non-tagged) frames received from an ISL trunk port are dropped.

- IEEE 802.1Q trunk port - Supports simultaneous tagged and untagged traffic. An 802.1Q trunk port is assigned a default Port VLAN ID (PVID), and all untagged traffic travels on the port default PVID. All untagged traffic and tagged traffic with a NULL VLAN ID are assumed to belong to the port default PVID. A packet with a VLAN ID equal to the outgoing port default PVID is sent untagged. All other traffic is sent with a VLAN tag.

Tunnel Ports

Tunnel ports are used in 802.1Q tunneling to segregate the traffic of customers in a service provider network from other customers who appear to be on the same VLAN. You configure an asymmetric link from a tunnel port on a service provider edge switch to an 802.1Q trunk port on the customer switch. Packets entering the tunnel port on the edge switch, already 802.1Q-tagged with the customer VLANs, are encapsulated with another layer of 802.1Q tag (called the metro tag) containing a VLAN ID unique in the service provider network, for each customer. The double-tagged packets go through the service-provider network keeping the original customer VLANs separate from those of other customers. At the outbound interface, also a tunnel port, the metro tag is removed, and the original VLAN numbers from the customer network are retrieved.

Tunnel ports cannot be trunk ports or access ports and must belong to a VLAN unique for each customer.

Note	Switch ports are configured using the **switchport** interface configuration command.

Access Port Configuration

This section discusses the configuration of Access Ports on the Catalyst 3550.

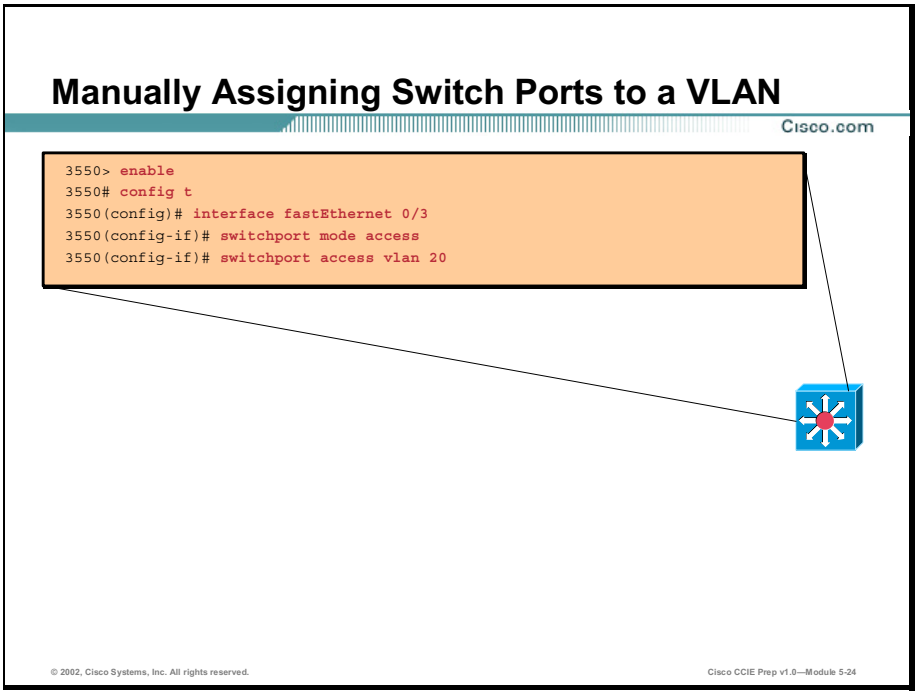

You can manually assign access ports to a VLAN without having VTP globally propagate VLAN configuration information.

Note	If you assign an interface to a VLAN that does not exist, a new VLAN is created.

Use the steps outlined in the following table to manually assign an access port to a VLAN.

Table 4-10: Assign Ports to a VLAN

Command	Purpose
3550(config)# **interface** *interface-id*	Enter the interface to be added to the VLAN.
3550(config-if)# **switchport mode** access	Define the VLAN membership mode for the port (Layer 2 access port).
3550(config-if)# **switchport access** vlan *vlan-id*	Assign the port to a VLAN. Valid VLAN IDs are 1 to 4094; do not enter leading zeros.

Note	To return an interface to its default configuration, use the **default interface** *interface-id* interface configuration command.

Trunk Port Configuration

This section discusses the configuration of Trunk Ports on the Catalyst 3550.

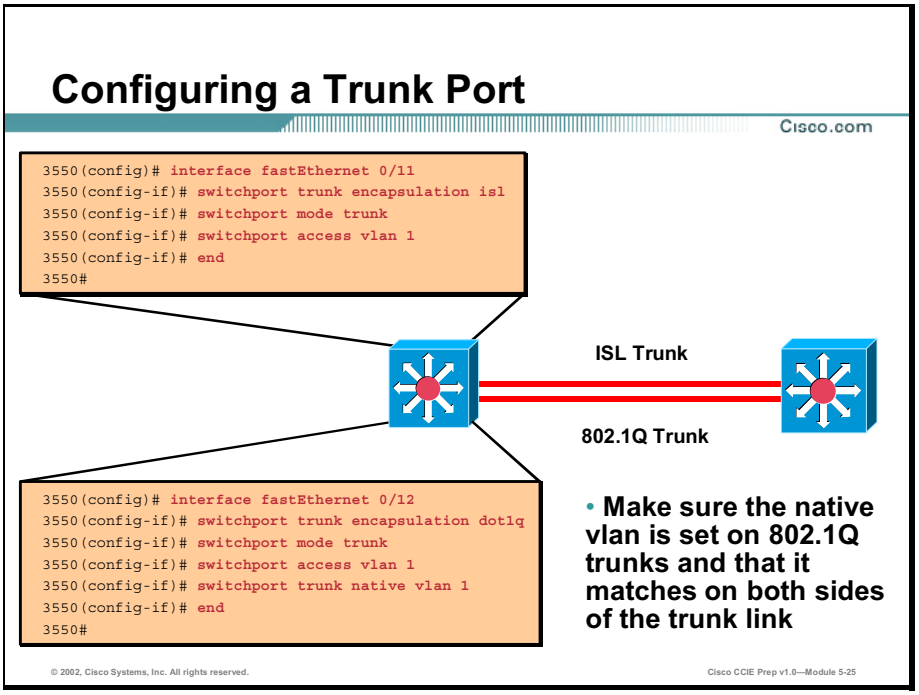

A trunk is a point-to-point link between one or more Ethernet switch interfaces and another networking device such as a router or a switch. Fast Ethernet and Gigabit Ethernet trunks carry the traffic of multiple VLANs over a single link, and you can extend the VLANs across an entire network.

Two trunking encapsulations are available on all Ethernet interfaces on the Catalyst 3550:

- Inter-Switch Link (ISL)—ISL is Cisco-proprietary trunking encapsulation.
- 802.1Q—802.1Q is industry-standard trunking encapsulation.

You can set an interface as trunking or nontrunking or to negotiate trunking with the neighboring interface. To autonegotiate trunking, the interfaces must be in the same VTP domain. Trunk negotiation is managed by the Dynamic Trunking Protocol (DTP), which is a Point-to-Point Protocol. DTP supports autonegotiation of both ISL and 802.1Q trunks.

Use the steps outlined in the following table to configure a port as an ISL or 802.1Q trunk port:

Table 4-11: Configure a Port

Command	Purpose
3550(config)# **interface** *interface-id*	Enter the interface configuration mode and the port to be configured for trunking. The default mode for Layer 2 interfaces is **switchport mode dynamic desirable**. If the neighboring interface supports trunking and is configured to allow trunking, the link is a Layer 2 trunk or, if the interface is currently in Layer 3 mode, it becomes a Layer 2 trunk when you enter the **switchport** interface configuration command.
3550(config-if)# **switchport trunk encapsulation** {**isl** | **dot1q** | **negotiate**}	Configure the port to support ISL or 802.1Q encapsulation or to negotiate (default) with the neighboring interface for encapsulation type. You must configure each end of the link with the same encapsulation type.
3550(config-if)# **switchport mode** {**dynamic** {**auto** | **desirable**} | **trunk**}	Configure the interface as a Layer 2 trunk (required only if the interface is currently a Layer 2 access port or tunnel port, or to specify the trunking mode). ■ **dynamic auto**—Set the interface to a trunk link if the neighboring interface is set to trunk or desirable mode. ■ **dynamic desirable**—Set the interface to a trunk link if the neighboring interface is set to trunk, desirable, or auto mode. ■ **trunk**—Set the interface in permanent trunking mode and negotiate to convert the link to a trunk link even if the neighboring interface is not a trunk interface.
3550(config-if)# **switchport access vlan** *vlan-id*	(Optional) Specify the default VLAN, which is used if the interface stops trunking.
3550(config-if)# **switchport trunk native vlan** *vlan-id*	Specify the native VLAN for 802.1Q trunks.

Note To reset all trunking characteristics of a trunking interface to the defaults, use the **no switchport trunk** interface configuration command. To disable trunking, use the **switchport mode access** interface configuration command to configure the port as a static-access port.

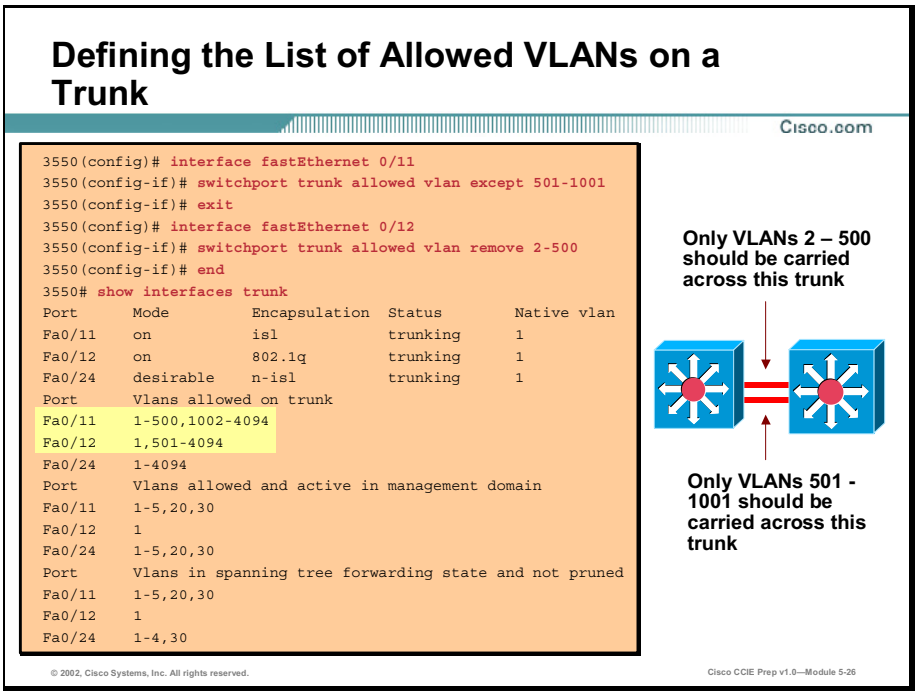

By default, a trunk port sends traffic to and receives traffic from all VLANs. All VLAN IDs, 1 to 4094, are allowed on each trunk. However, you can remove VLANs from the allowed list, preventing traffic from those VLANs from passing over the trunk. To restrict the traffic a trunk carries, use the **switchport trunk allowed vlan remove** *vlan-list* interface configuration command to remove specific VLANs from the allowed list.

A trunk port can become a member of a VLAN if the VLAN is enabled, if VTP knows of the VLAN, and if the VLAN is in the allowed list for the port. When VTP detects a newly enabled VLAN and the VLAN is in the allowed list for a trunk port, the trunk port automatically becomes a member of the enabled VLAN. When VTP detects a new VLAN and the VLAN is not in the allowed list for a trunk port, the trunk port does not become a member of the new VLAN.

Use the steps outlined in the following table to restrict the VLANs that are carried on a trunk port:

Table 4-12: Restrict VLANs

Command	Purpose
3550(config-if)# **switchport trunk allowed vlan** {**add** | **all** | **except** | **remove**} *vlan-list*	Configure the list of VLANs allowed on the trunk. The *vlan-list* parameter is either a single VLAN number from 1 to 4094 or a range of VLANs described by two VLAN numbers, the lower one first, separated by a hyphen. Do not enter any spaces between comma-separated VLAN parameters or in hyphen-specified ranges. All VLANs are allowed by default. You cannot remove any of the default VLANs (1 or 1002 to 1005) from a trunk.

Note	To return to the default allowed VLAN list of all VLANs, use the **no switchport trunk allowed vlan** interface configuration command.

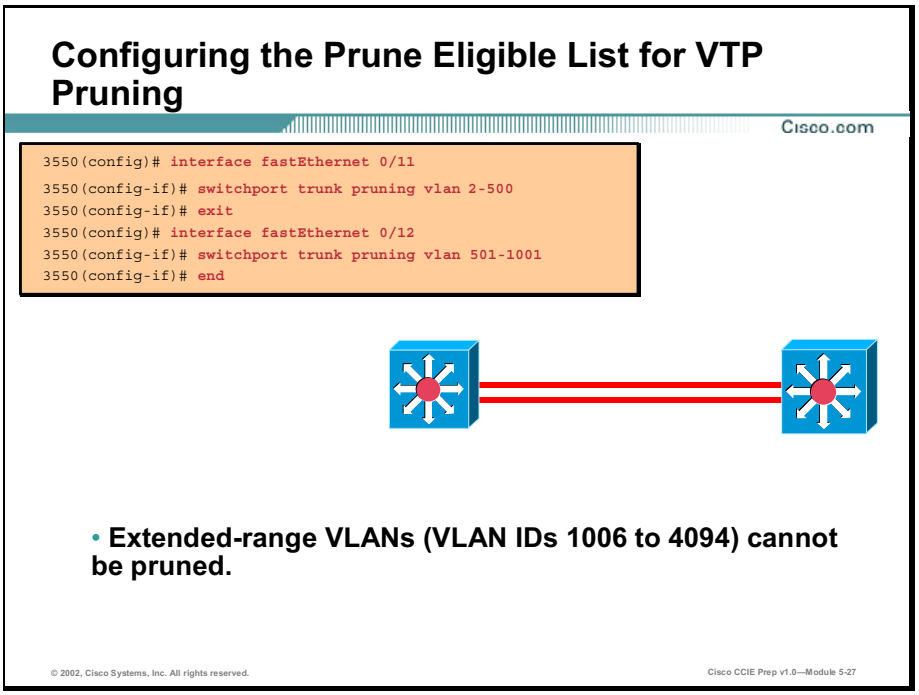

Pruning increases available bandwidth by restricting flooded traffic to those trunk links that the traffic must use to access the destination devices. You can only enable VTP pruning on a switch in VTP server mode. Only VLANs included in the pruning-eligible list can be pruned. By default, VLANs 2 through 1001 are pruning eligible on trunk ports. The pruning-eligible list applies only to trunk ports. Each trunk port has its own eligibility list. VTP pruning must be enabled for this procedure to take effect.

Use the steps outlined in the following table to remove VLANs from the pruning-eligible list on a trunk port:

Table 4-13: Remove VLANs from Pruning

Command	Purpose
3550(config-if)# **switchport trunk pruning vlan** {**add** | **except** | **none** | **remove**} *vlan-list* [,*vlan*[,*vlan*[,,,]]	Configure the list of VLANs allowed to be pruned from the trunk. Separate nonconsecutive VLAN IDs with a comma and no spaces; use a hyphen to designate a range of IDs. Valid IDs are from 2 to 1001**. Extended-range VLANs (VLAN IDs 1006 to 4094) cannot be pruned.** VLANs that are pruning-ineligible receive flooded traffic. The default list of VLANs allowed to be pruned contains VLANs 2 to 1001.

Note	To return to the default pruning-eligible list of all VLANs, use the **no switchport trunk pruning vlan** interface configuration command.

Tunnel Port Configuration

This section discusses the configuration of Tunnel Ports on the Catalyst 3550. Tunneling is a feature designed for service providers who carry traffic of multiple customers across their networks and are required to maintain the VLAN and Layer 2 protocol configurations of each customer without impacting the traffic of other customers. The Catalyst 3550 switch supports 802.1Q tunneling and Layer 2 protocol tunneling.

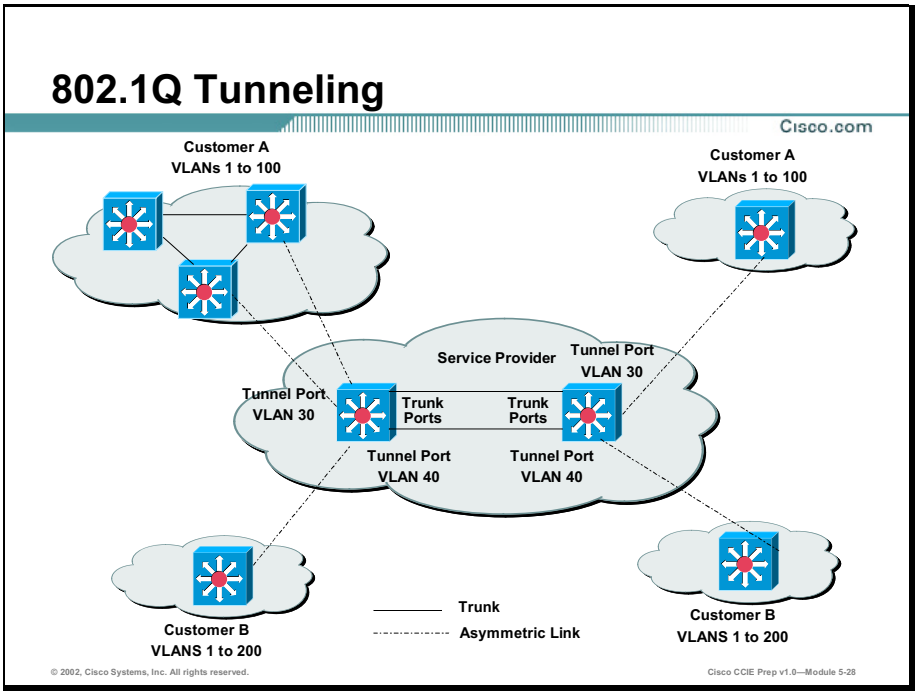

Business customers of service providers often have specific requirements for VLAN IDs and the number of VLANs to be supported. The VLAN ranges required by different customers in the same service-provider network might overlap, and traffic of customers through the infrastructure might be mixed. Assigning a unique range of VLAN IDs to each customer would restrict customer configurations and could easily exceed the VLAN limit of 4096 of the 802.1Q specification.

Using the 802.1Q tunneling feature, service providers can use a single VLAN to support customers who have multiple VLANs. Customer VLAN IDs are preserved and traffic from different customers is segregated within the service-provider infrastructure even when they appear to be on the same VLAN. The 802.1Q tunneling expands VLAN space by using a VLAN-in-VLAN hierarchy and tagging the tagged packets. A port configured to support 802.1Q tunneling is called a tunnel port. When you configure tunneling, you assign a tunnel port to a VLAN that is dedicated to tunneling. Each customer requires a separate VLAN, but that VLAN supports all of the customer's VLANs.

Customer traffic tagged in the normal way with appropriate VLAN IDs come from an 802.1Q trunk port on the customer device and into a tunnel port on the service-provider edge switch. The link between the customer device and the edge switch is an asymmetric link because one

end is configured as an 802.1Q trunk port and the other end is configured as a tunnel port. You assign the tunnel port interface to an access VLAN ID unique to each customer.

Packets coming from the customer trunk port into the tunnel port on the service-provider edge switch are normally 802.1Q-tagged with appropriate VLAN ID. The tagged packets remain intact inside the switch and, when they exit the trunk port into the service provider network, are encapsulated with another layer of an 802.1Q tag (called the *metro tag*) that contains the VLAN ID unique to the customer. The original 802.1Q tag from the customer is preserved in the encapsulated packet. Therefore, packets entering the service-provider infrastructure are double-tagged, with the outer tag containing the customer's access VLAN ID, and the inner VLAN ID being the VLAN of the incoming traffic.

When the double-tagged packet enters another trunk port in a service provider core switch, the outer tag is stripped as the packet is processed inside the switch. When the packet exits another trunk port on the same core switch, the same metro tag is again added to the packet.

When the packet enters the trunk port of the service-provider egress switch, the outer tag is again stripped as the packet is processed internally on the switch. However, the metro tag is not added when it is sent out the tunnel port on the edge switch into the customer network, and the packet is sent as a normal 802.1Q-tagged frame to preserve the original VLAN numbers in the customer network.

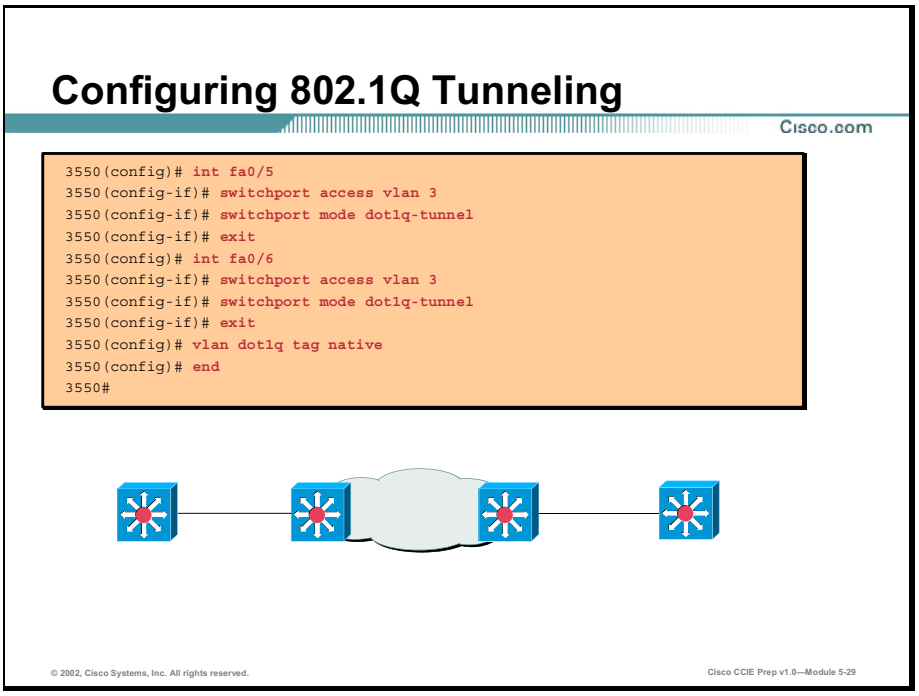

When you configure 802.1Q tunneling, you should always use asymmetrical links for traffic going in or out of a tunnel and dedicate one VLAN for each tunnel. You should also be aware of configuration requirements for native VLANs and maximum transmission units (MTU).

Native VLANs

When configuring 802.1Q tunneling on an edge switch, you must use 802.1Q trunk ports for sending out packets into the service-provider network. However, packets going through the core of the service-provider network might be carried through 802.1Q trunks, ISL trunks, or nontrunking links. When 802.1Q trunks are used in these core switches, the native VLANs of the 802.1Q trunks must not match any native VLAN of the nontrunking (tunneling) port on the same switch because traffic on the native VLAN would not be tagged on the 802.1Q transmitting trunk port.

These are some ways to solve this problem:

- Use ISL trunks between core switches in the service-provider network. Although customer interfaces connected to edge switches must be 802.1Q trunks, we recommend using ISL trunks for connecting switches in the core layer.

- Configure the edge switch so that all packets going out an 802.1Q trunk, including the native VLAN, are tagged by using the **vlan dot1q tag native** global configuration command. If the switch is configured to tag native VLAN packets on all 802.1Q trunks, the switch accepts untagged packets, but sends only tagged packets.

- Ensure that the native VLAN ID on the edge switch trunk port is not within the customer VLAN range. For example, if the trunk port carries traffic of VLANs 100 to 200, assign the native VLAN a number outside that range.

System MTU

The default system MTU for traffic on the Catalyst 3550 switch is 1500 bytes. You can configure the switch to support frames larger than 1500 bytes by using the **system mtu** global configuration command. Because the 802.1Q tunneling feature increases the frame size by 4 bytes when the metro tag is added, you must configure all switches in the service-provider network to be able to process maximum frames by increasing the switch system MTU size to at least 1504 bytes. The maximum allowable system MTU for Catalyst 3550 Gigabit Ethernet switches is 2000 bytes; the maximum system MTU for Fast Ethernet switches is 1546 bytes.

Use the steps outlined in the following table to configure a switch port on the Catalyst 3550 as an 802.1Q tunnel port:

Table 4-14: Configure a Switch Port

Command	Purpose
3550(config)# **interface** *interface-id*	Enter interface configuration mode and the interface to be configured as a tunnel port. This should be the edge port in the service-provider network that connects to the customer switch. Valid interfaces include physical interfaces and port-channel logical interfaces (port channels 1 to 64).
3550(config-if)# **switchport access vlan** *vlan-id*	Specify the default VLAN, which is used if the interface stops trunking. This is VLAN ID specific to the particular customer.
3550(config-if)# **switchport mode dot1q-tunnel**	Set the interface as an 802.1Q tunnel port.
3550(config)# **vlan dot1q tag native**	(Optional) Set the switch to enable tagging of native VLAN packets on all 802.1Q trunk ports. When not set, if a customer VLAN ID is the same as the native VLAN, the trunk port does not apply a metro tag, and packets might be sent to the wrong destination.

Use the **no switchport mode dot1q-tunnel** interface configuration command to return the port to the default state of dynamic desirable. Use the **no vlan dot1q tag native** global configuration command to disable tagging of native VLAN packets.

Although 802.1Q tunneling works well for Layer 2 packet switching, there are incompatibilities with some Layer 2 features and with Layer 3 switching.

- A tunnel port cannot be a routed port.
- IP routing is not supported on a VLAN that includes 802.1Q ports. Packets received from a tunnel port are forwarded based only on Layer 2 information. If routing is enabled on the switch virtual interface (SVI) that includes tunnel ports, untagged IP packets received from the tunnel port are recognized and routed by the switch. This allows the customer to access the internet through its native VLAN. If this access is not required, you should not configure SVIs on VLANs that include tunnel ports.
- Fallback bridging is not supported on tunnel ports. Because all 802.1Q-tagged packets received from a tunnel port are treated as non-IP packets, if fallback bridging is enabled on VLANs that have tunnel ports configured, IP packets would be improperly bridged across VLANs. Therefore, you must *not* enable fallback bridging on VLANs with tunnel ports.

- Tunnel ports do not support IP access control lists (ACLs).
- Layer 3 quality of service (QoS) ACLs and other QoS features related to Layer 3 information are not supported on tunnel ports. MAC-based QoS is supported on tunnel ports.
- EtherChannel port groups are compatible with tunnel ports as long as the 802.1Q configuration is consistent within an EtherChannel port group.
- PAgP and UDLD are supported on 802.1Q tunnel ports.
- Dynamic Trunking Protocol (DTP) is not compatible with 802.1Q tunneling because you must manually configure asymmetric links with tunnel ports and trunk ports.
- Loopback detection is supported on 802.1Q tunnel ports.
- When a port is configured as an 802.1Q tunnel port, spanning tree bridge protocol data unit (BPDU) filtering is automatically enabled on the interface. Cisco Discovery Protocol (CDP) is automatically disabled on the interface.

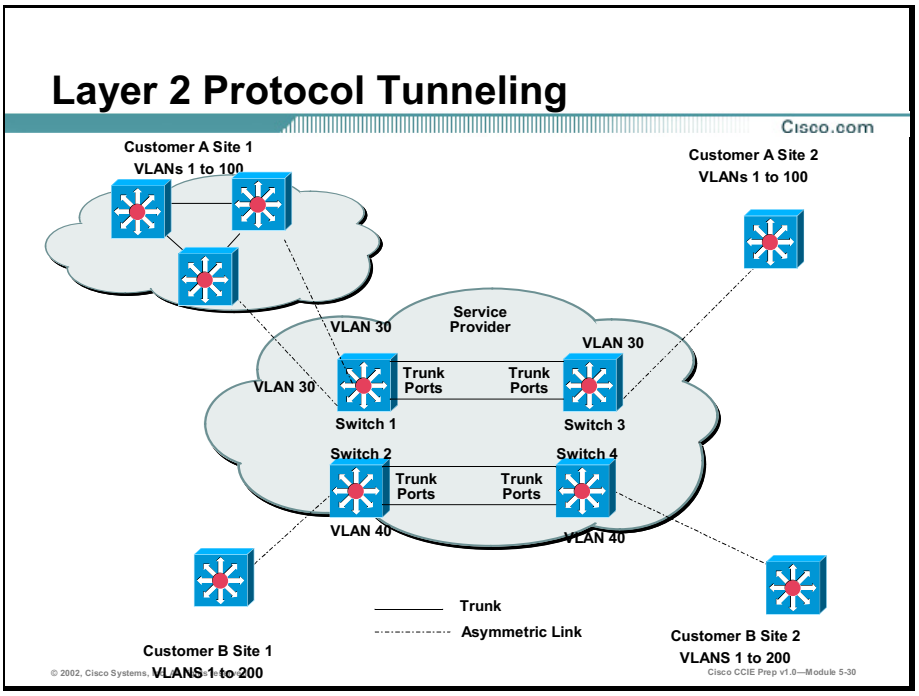

Customers that have different sites connected across a service-provider network and want to scale this topology into one large layer 2 domain need to run various Layer 2 protocols between sites. For example, STP must run properly, and every VLAN should build a proper spanning tree that includes the local site and all remote sites across the service-provider infrastructure. Cisco Discovery Protocol (CDP) must discover neighboring Cisco devices from local and remote sites. VLAN Trunking Protocol (VTP) must provide consistent VLAN configuration throughout all sites in the customer network.

When protocol tunneling is enabled, edge switches on the inbound side of the service-provider infrastructure encapsulate Layer 2 protocol packets with a special MAC address and send them across the service-provider network. Core switches in the network do not process these packets, but forward them as normal packets. Layer 2 protocol data units (PDUs) for CDP, STP, or VTP cross the service-provider infrastructure and are delivered to customer switches on the outbound side of the service-provider network. Identical packets are received by all customer ports on the same VLANs with the following results:

- Switches at each of a customer's sites are able to properly run STP and every VLAN can build a correct spanning tree based on parameters from all sites and not just from the local site.
- CDP discovers and shows information about the other Cisco devices connected through the service-provider network.
- VTP provides consistent VLAN configuration throughout the customer network, propagating through the service provider to all switches.

Layer 2 protocol tunneling can be used independently or to enhance 802.1Q tunneling. If protocol tunneling is not enabled on 802.1Q tunneling ports, remote switches at the receiving end of the service-provider network do not receive the PDUs and cannot properly run STP,

CDP, and VTP. When protocol tunneling *is* enabled, Layer 2 protocols within each customer's network are totally separate from those running within the service-provider network. Customer switches at different sites that send traffic through the service-provider network with 802.1Q tunneling achieve complete knowledge of the customer's VLAN. If 802.1Q tunneling is not used, you can still enable Layer 2 protocol tunneling by connecting to the customer switch through access ports and enabling tunneling on the service-provider access port.

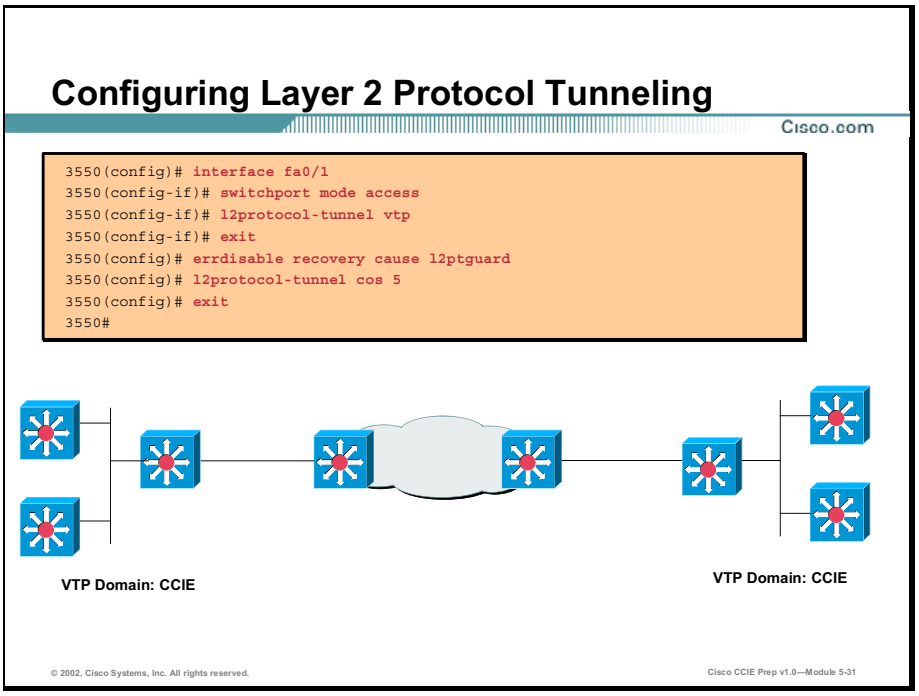

You enable Layer 2 protocol tunneling (by protocol) on the access ports or tunnel ports that are connected to the customer in the edge switches of the service-provider network. Edge-switch tunnel ports are connected to customer 802.1Q trunk ports; edge-switch access ports are connected to customer access ports. The Catalyst 3550 switch supports Layer 2 protocol tunneling for CDP, STP, and VTP. The edge switches connected to the customer switch perform the tunneling process.

When the Layer 2 PDUs that entered the inbound edge switch through the tunnel or access port exit the switch through the trunk port into the service-provider network, the switch overwrites the customer PDU-destination MAC address with a well-known Cisco proprietary multicast address (01-00-0c-cd-cd-d0). If 802.1Q tunneling is enabled, packets are also double-tagged; the outer tag is the customer metro tag and the inner tag is the customer VLAN tag. The core switches ignore the inner tags and forward the packet to all trunk ports in the same metro VLAN. The edge switches on the outbound side restore the proper Layer 2 protocol and MAC address information and forward the packets to all tunnel or access ports in the same metro VLAN. Therefore, the Layer 2 PDUs are kept intact and delivered across the service-provider infrastructure to the other side of the customer network.

Use the steps outlined in the following table to configure a port for Layer 2 protocol tunneling:

Table 4-15: Configure a Port for Layer 2 Protocol Tunneling

Command	Purpose
3550(config)# **interface** *interface-id*	Enter the interface configuration mode and the interface to be configured as a tunnel port. This should be the edge port in the service-provider network that connects to the customer switch. Valid interfaces include physical interfaces and port-channel logical interfaces (port channels 1 to 64).
3550(config-if)# **switchport mode access** or 3550(config-if)# **switchport mode dot1q-tunnel**	Configure the interface as an access port or an 802.1Q tunnel port.
3550(config-if)# **l2protocol-tunnel** {**cdp** | **vtp** | **stp**}	Enable protocol tunneling for the desired protocol.
3550(config-if)# **l2protocol-tunnel shutdown-threshold** {**cdp** | **vtp** | **stp**} *value*	(Optional) Configure the threshold in packets per second to be received for encapsulation and transmitted before the interface shuts down. The threshold is based on the combined (linear) sum of the rate at which the specific L2 protocol packets are received and the rate at which the specific L2 protocol packets are transmitted on the port. The port is disabled if the configured threshold is exceeded. The range is 1 to 4096. The default is to have no threshold configured.
3550(config)# **errdisable recovery cause l2ptguard**	(Optional) Configure the recovery mechanism from a Layer 2 maximum rate error so that the interface can be brought out of the disabled state and allowed to try again. You can also set the time interval. Errdisable recovery is disabled by default; when enabled, the default time interval is 300 seconds.
3550(config)# **l2protocol-tunnel cos** *value*	(Optional) Configure the CoS value for all tunneled Layer 2 PDUs. The range is 0 to 7; the default is the default COS value for the interface. If none is configured, the default is 5.

These are some configuration guidelines and operating characteristics of Layer 2 protocol tunneling:

- The switch supports tunneling of CDP, STP, including multiple STP (MSTP), and VTP protocols. Protocol tunneling is disabled by default but can be enabled for the individual protocols on 802.1Q tunnel ports or on access ports.
- Tunneling is not supported on trunk ports.If you enter the **l2protocol-tunnel** interface configuration command on a trunk port, the command is accepted, but Layer 2 tunneling does not take affect unless you change the port to a tunnel port or access port.
- EtherChannel port groups are compatible with tunnel ports as long as the 802.1Q configuration is consistent within an EtherChannel port group.
- If an encapsulated PDU (with the proprietary destination MAC address) is received from a tunnel port or access port with Layer 2 tunneling enabled, the tunnel port is shut down to prevent loops. The port also shuts down when a configured shutdown threshold for the protocol is reached. You can manually re-enable the port (by issuing a **shutdown**, **no shutdown** command sequence) or if errdisable recovery is enabled, the operation is retried after a specified time interval.

- Only decapsulated PDUs are forwarded to the customer network. The spanning tree instance running on the service-provider network does not forward BPDUs to tunnel ports. No CDP packets are forwarded from tunnel ports.

- When protocol tunneling is enabled on an interface, you can set a per protocol, per port, shutdown threshold for the PDUs generated by the customer network. If the limit is exceeded, the port is shut down. You can also rate-limit BPDUs by using QoS ACLs and policy maps on a tunnel port.

- Because tunneled PDUs (especially STP BPDUs) must be delivered to all remote sites for the customer virtual network to operate properly, you can give PDUs higher priority within the service-provider network than data packets received from the same tunnel port. By default, the PDUs use the same CoS value as data packets.

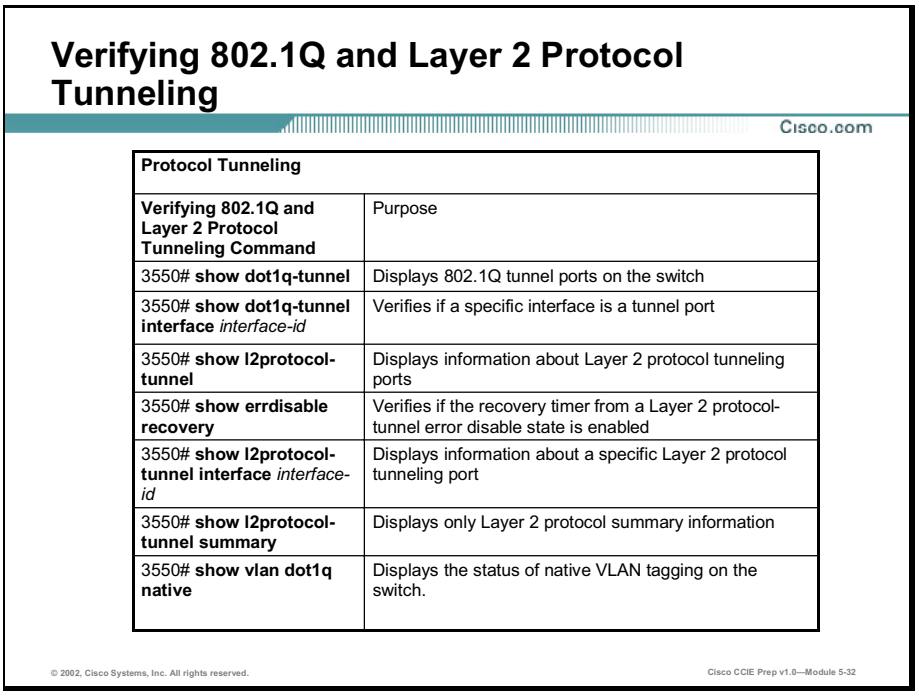

The following table lists the commands used for verifying 802.1Q and Layer 2 Protocol Tunneling:

Table 4-16: Protocol Tunneling

Verifying 802.1Q and Layer 2 Protocol Tunneling Command	Purpose
3550# **show dot1q-tunnel**	Displays 802.1Q tunnel ports on the switch.
3550# **show dot1q-tunnel interface** *interface-id*	Verifys if a specific interface is a tunnel port.
3550# **show l2protocol-tunnel**	Displays information about Layer 2 protocol tunneling ports.
3550# **show errdisable recovery**	Verifys if the recovery timer from a Layer 2 protocol-tunnel error disable state is enabled.
3550# **show l2protocol-tunnel interface** *interface-id*	Displays information about a specific Layer 2 protocol tunneling port.
3550# **show l2protocol-tunnel summary**	Displays only Layer 2 protocol summary information.
3550# **show vlan dot1q native**	Displays the status of native VLAN tagging on the switch.

Layer 3 Interfaces

The Catalyst 3550 supports two different types of Layer 3 interfaces: Router Ports, which are physical interfaces that act just like a physical interface on a Cisco IOS router, and Switched Virtual Interfaces (SVI), which are virtual VLAN interfaces used for InterVLAN routing, similar to the VLAN interfaces on the MSFC of the Catalyst 6500 series.

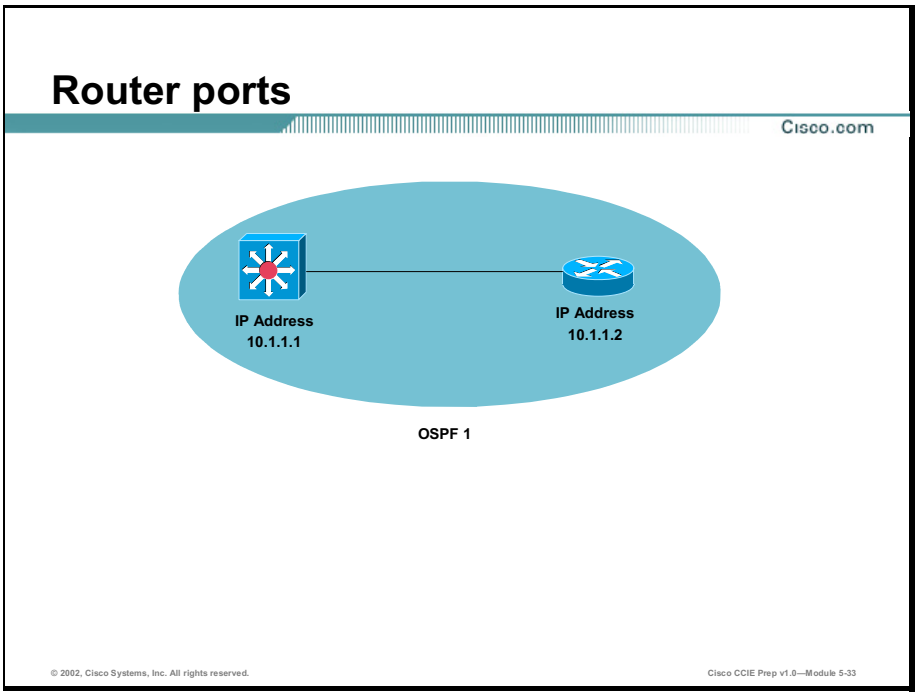

A routed port is a physical port that acts like a port on a router; it does not have to be connected to a router. A routed port is not associated with a particular VLAN, as is an access port. A routed port behaves like a regular router interface, except that it does not support VLAN subinterfaces. Routed ports can be configured with a Layer 3 routing protocol.

Configure routed ports by putting the interface into Layer 3 mode with the **no switchport** interface configuration command. Then assign an IP address to the port, enable routing, and assign routing protocol characteristics by using the **ip routing** and **router** *protocol* global configuration commands.

Note	Entering a **no switchport** interface configuration command shuts the interface down and then re-enables it, which might generate messages on the device to which the interface is connected. Furthermore, when you use this command to put the interface into Layer 3 mode, you are deleting any Layer 2 characteristics configured on the interface.

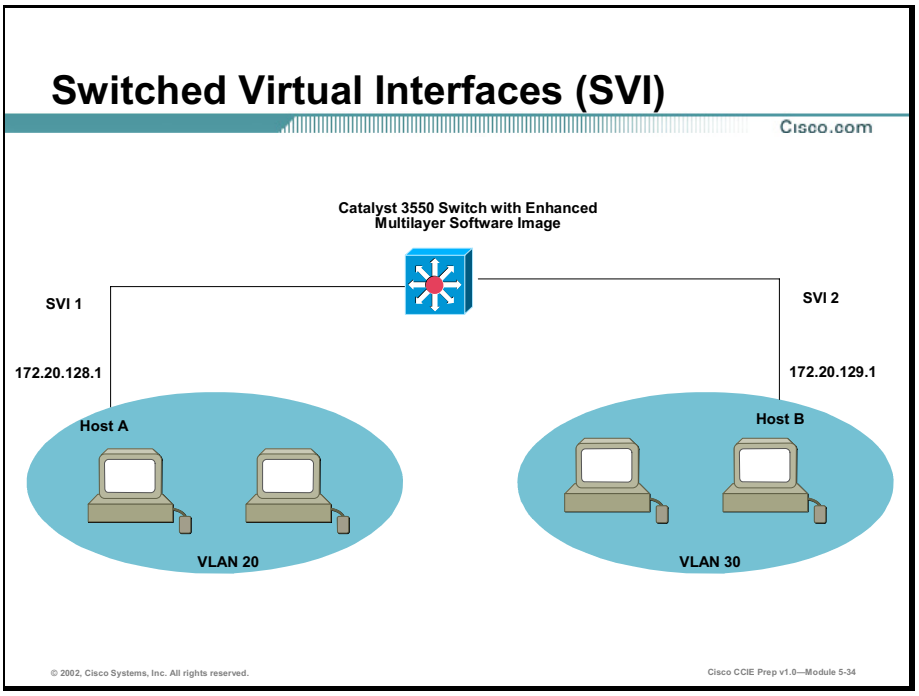

A switch virtual interface (SVI) represents a VLAN of switch ports as one interface to the routing or bridging function in the system. Only one SVI can be associated with a VLAN, but you need to configure an SVI for a VLAN only when you wish to route between VLANs, fallback-bridge nonroutable protocols between VLANs, or to provide IP host connectivity to the switch. By default, an SVI is created for the default VLAN (VLAN 1) to permit remote switch administration. Additional SVIs must be explicitly configured. In Layer 2 mode, SVIs provide IP host connectivity only to the system; in Layer 3 mode, you can configure routing across SVIs.

SVIs are created the first time that you enter the **vlan** interface configuration command for a VLAN interface. The VLAN corresponds to the VLAN tag associated with data frames on an ISL or 802.1Q encapsulated trunk or the VLAN ID configured for an access port. Configure a VLAN interface for each VLAN for which you want to route traffic, and assign it an IP address.

Note When you create an SVI, it does not become active until it is associated with a physical port.

Routed ports and SVIs support routing protocols (including Multicast routing) and bridging configurations. The process of configuring IP addresses and routing protocols on the Catalyst 3550 is the same as any IOS-based device (Cisco router). Many of the IOS commands learned in this course also apply to the Catalyst 3550. All Layer 3 interfaces require an IP address to route traffic.

The number of routed ports and SVIs that you can configure is not limited by software; however, the interrelationship between this number and the number of other features being configured might have an impact on CPU utilization because of hardware limitations.

General Interface Commands

The following commands are applicable to all interfaces; logical or physical, Layer 2 or Layer 3, on the Catalyst 3550.

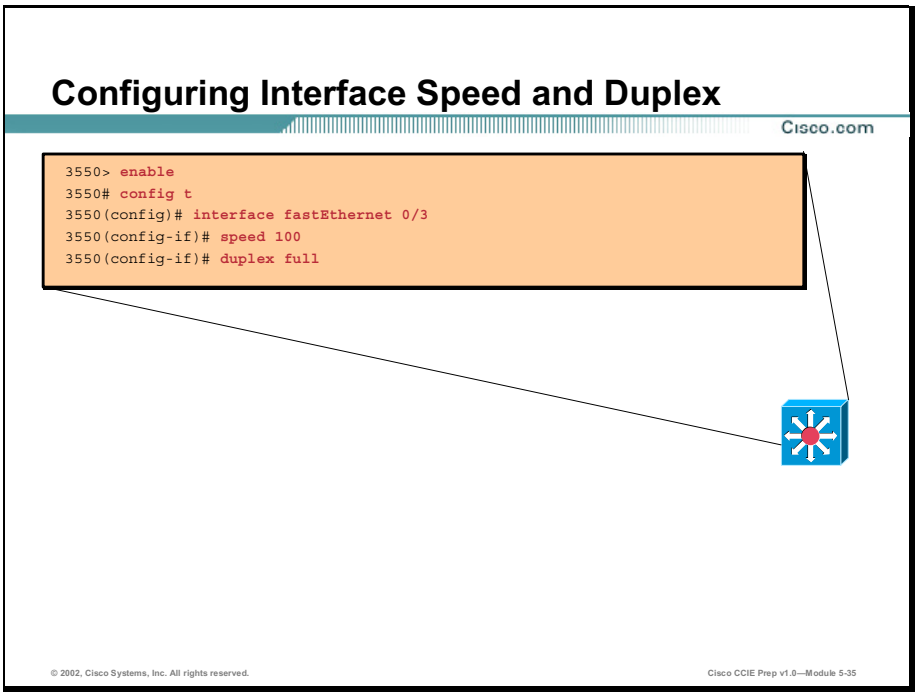

Ethernet interfaces on the Catalyst 3550 switch operate in 10, 100, or 1000 Mbps and in either full or half duplex mode. In full-duplex mode, two stations can send and receive at the same time. When packets can flow in both directions simultaneously, effective Ethernet bandwidth doubles to 20 Mbps for Ethernet interfaces, to 200 Mbps for Fast Ethernet interfaces, and to 2 Gbps for Gigabit interfaces. Full-duplex communication is often an effective solution to collisions, which are major constrictions in Ethernet networks. Normally, Ethernet ports operate in half-duplex mode, which means that stations can either receive or send.

You can configure interface speed on Fast Ethernet (10/100-Mbps) and Gigabit Ethernet (10/100/1000-Mbps) interfaces; you cannot configure speed on Gigabit Interface Converter (GBIC) interfaces. You can configure duplex mode on any Fast Ethernet or Gigabit Ethernet interfaces that are not set to autonegotiate; you cannot configure duplex mode on GBIC interfaces.

Use the steps outlined in the following table to set the speed and duplex mode for a 10/100/1000 Ethernet interface:

Table 4-17: Set Speed and Duplex Mode

Command	Purpose
3550(config)# **interface** *interface-id*	Enters interface configuration mode.
3550(config-if)# **speed** {**10** | **100** | **1000** | **auto** | **nonegotiate**}	Enters the appropriate speed parameter for the interface. Other options are **auto** or **nonegotiate**. Note: The **1000** keyword is available only for 10/100/1000 Mbps ports. 100BASE-FX ports operate only at 100 Mbps. GBIC module ports operate only at 1000 Mbps. The **nonegotiate** keyword is available only for 1000BASE-SX, -LX, and -ZX GBIC ports.
3550(config-if)# **duplex** {**auto** | **full** | **half**}	Enters the duplex parameter for the interface. Note: 100BASE-FX ports operate only in full-duplex mode.

Note Use the **no speed** and **no duplex** interface configuration commands to return the interface to the default speed and duplex settings (autonegotiate). To return all interface settings to the defaults, use the **default interface** *interface-id* interface configuration command.

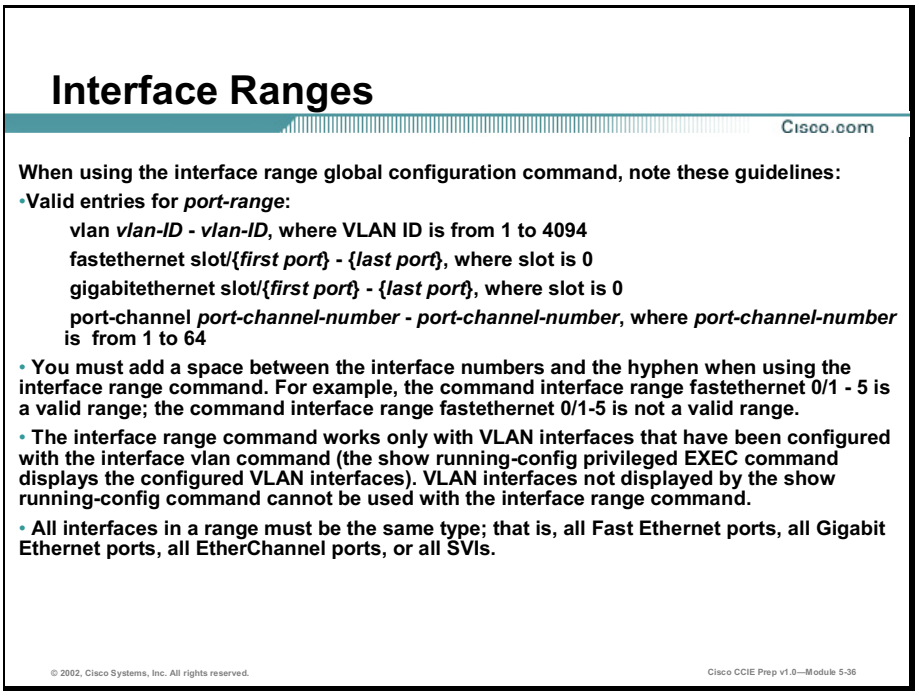

You can use the **interface range** global configuration command to configure multiple interfaces with the same configuration parameters. When you enter the **interface range** configuration mode, all command parameters that you enter are attributed to all interfaces within that range until you exit this mode. The **interface range** command is extremely useful when assigning access ports to a VLAN. Without the interface range command you must go into interface configuration on each individual access port.

When using the **interface range** global configuration command, note these guidelines:

- Valid entries for *port-range*:
  - **vlan** *vlan-ID* - *vlan-ID*, where VLAN ID is from 1 to 4094
  - **fastethernet** slot/{*first port*} - {*last port*}, where slot is **0**
  - **gigabitethernet** slot/{*first port*} - {*last port*}, where slot is **0**
  - **port-channel** *port-channel-number - port-channel-number*, where *port-channel-number* is from 1 to 64
- You must add a space between the interface numbers and the hyphen when using the **interface range** command. For example, the command **interface range fastethernet 0/1 - 5** is a valid range; the command **interface range fastethernet 0/1-5** is not a valid range.
- The **interface range** command works only with VLAN interfaces that have been configured with the **interface vlan** command (the **show running-config** privileged EXEC

command displays the configured VLAN interfaces). VLAN interfaces not displayed by the **show running-config** command cannot be used with the **interface range** command.

- All interfaces in a range must be the same type; that is, all Fast Ethernet ports, all Gigabit Ethernet ports, all EtherChannel ports, or all SVIs.

Use the steps outlined in the following table to configure a range of interfaces with the same parameters:

Table 4-18: Range of Interfaces

Command	Purpose
3550(config)# **interface range** {*port-range*}	Enter interface range configuration mode by entering the range of interfaces (VLANs or physical ports) to be configured. ■ You can use the **interface range** command to configure up to five port ranges or a previously defined macro. ■ Each comma-separated *port-range* must consist of the same port type. You do not need to enter spaces before or after the comma. When you define a range, the space between the first port and the hyphen is required.

You can now use the normal configuration commands to apply the configuration parameters to all interfaces in the range.

You can create an interface range macro to automatically select a range of interfaces for configuration. Before you can use the **macro** keyword in the **interface range macro** global configuration command string, you must use the **define interface-range** global configuration command to define the macro.

Use the steps outlined in the following table to define an interface range macro:

Table 4-19: Interface Range Macro

Command	Purpose
3550(config)# **define interface-range** *macro_name interface-range*	Define the interface-range macro, and save it in NVRAM. ■ The *macro_name* is a 32-character maximum character string. ■ A macro can contain up to five comma-separated interface ranges. You do not need to enter spaces before or after the comma. ■ Each *interface-range* must consist of the same port type. All interfaces in a range must be the same type; that is, all Fast Ethernet ports, all Gigabit Ethernet ports, all EtherChannel ports, or all VLANs, but you can combine multiple interface types in a macro.
3550(config)# **interface range macro** *macro_name*	Select the interface range to be configured using the values saved in the interface-range macro called *macro_name.* You can now use the normal configuration commands to apply the configuration to all interfaces in the defined macro.

Note Use the **no define interface-range** *macro_name* global configuration command to delete a macro.

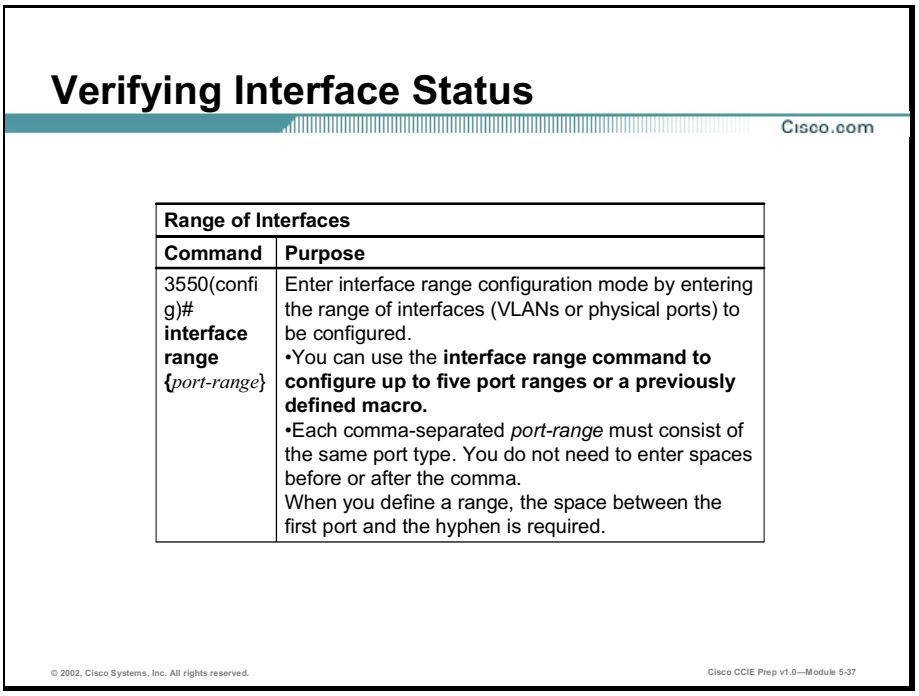

The table below lists the various Cisco IOS commands that can be used to verify the status of the interfaces on the Catalyst 3550.

Table 4-20: IOS Commands

Command	Purpose
3550# **show interfaces** [*interface-id*]	Displays the status and configuration of all interfaces or a specific interface.
3550# **show interfaces** *interface-id* **status** [**err-disabled**]	Displays interface status or a list of interfaces in error-disabled state.
3550# **show interfaces** [*interface-id*] **switchport**	Displays administrative and operational status of switch ports. You can use this command to determine if a port is in routing or switching mode.
3550# **show interfaces** [*interface-id*] **description**	Displays the description configured on an interface or all interfaces and their status.
3550# **show ip interface** [*interface-id*]	Displays the usability status of all interfaces configured for IP or the specified interface.
3550# **show ip interface brief**	Displays a brief summary of all IP interfaces.
3550# **show running-config interface** [*interface-id*]	Displays the running configuration of a particular interface.
3550# **show interfaces** [*interface-id*]	Displays the status and configuration of all interfaces or a specific interface.

EtherChannel

This section describes how to configure EtherChannel on Layer 2 and Layer 3 interfaces. EtherChannel provides fault-tolerant high-speed links between switches, routers, and servers.

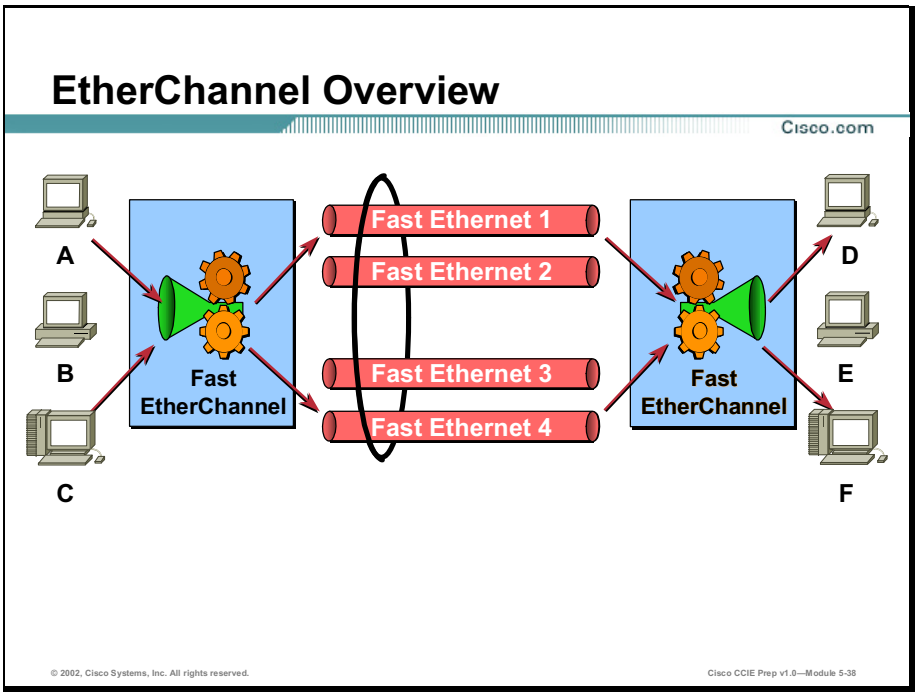

An EtherChannel consists of individual Fast Ethernet or Gigabit Ethernet links bundled into a single logical link. The EtherChannel provides full-duplex bandwidth up to 800 Mbps (Fast EtherChannel) or 8 Gbps (Gigabit EtherChannel) between your switch and another switch or host. Each EtherChannel can consist of up to eight compatibly configured Ethernet interfaces. All interfaces in each EtherChannel must be the same speed, and all must be configured as either Layer 2 or Layer 3 interfaces.

Etherchannel can be deployed anywhere in the network where bottlenecks are likely to occur. A common example is deploying Etherchannel between the wiring closets and the data center to increase bandwidth to support the aggregate bandwidth requirements of clients. EtherChannel also provides automatic recovery for the loss of a link by redistributing the load across the remaining links. If a link fails, EtherChannel redirects traffic from the failed link to the remaining links in the channel without intervention.

On the Catalyst 3550 Etherchannels can be created on both Layer 2 (switch ports) and Layer 3 (router ports) interfaces. The configuration differs between the however both configurations involve logical interfaces.

- With Layer 3 interfaces, you manually create the logical interface by using the **interface port-channel** global configuration command.
- With Layer 2 interfaces, the logical interface is dynamically created.

- With both Layer 3 and 2 interfaces, you manually assign an interface to the EtherChannel by using the **channel-group** interface configuration command. This command binds the physical and logical ports together.

Each EtherChannel has a logical port-channel interface numbered from 1 to 64. The channel groups are also numbered from 1 to 64.

After you configure an EtherChannel, configuration changes applied to the port-channel interface apply to all the physical interfaces assigned to the port-channel interface. Configuration changes applied to the physical interface affect only the interface where you apply the configuration. To change the parameters of all ports in an EtherChannel, apply configuration commands to the port-channel interface.

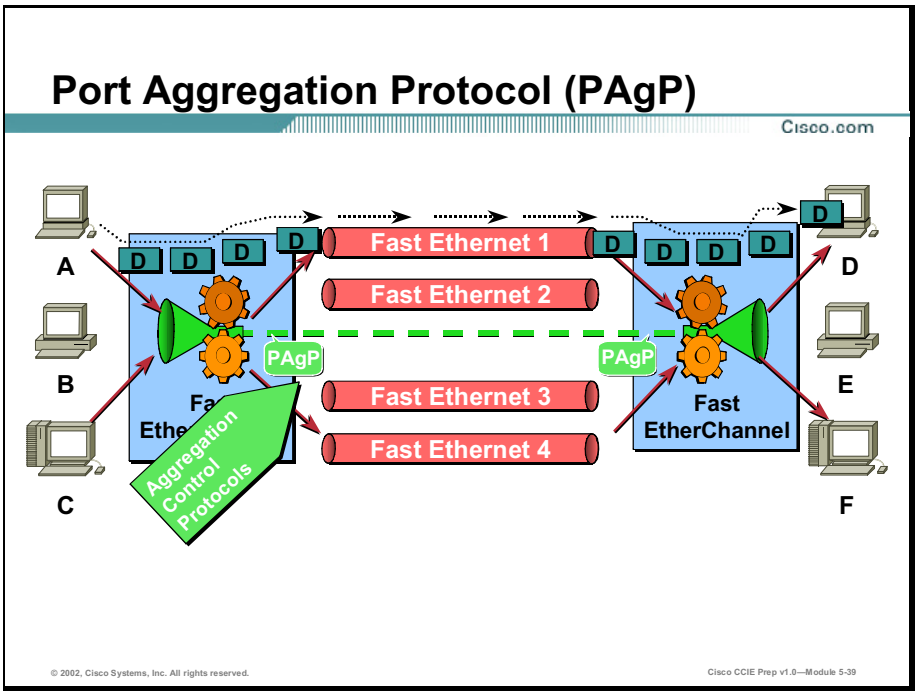

The Port Aggregation Protocol (PAgP) facilitates the automatic creation of EtherChannels by exchanging packets between Ethernet interfaces. By using PAgP, the switch learns the identity of partners capable of supporting PAgP and learns the capabilities of each interface. It then dynamically groups similarly configured interfaces into a single logical link (channel or aggregate port); these interfaces are grouped based on hardware, administrative, and port parameter constraints. For example, PAgP groups the interfaces with the same speed, duplex mode, native VLAN, VLAN range, and trunking status and type. After grouping the links into an EtherChannel, PAgP adds the group to the spanning tree as a single switch port.

PAgP Modes

The table below shows the user-configurable EtherChannel modes for the **channel-group** interface configuration command: **on**, **auto**, and **desirable**. Switch interfaces exchange PAgP packets only with partner interfaces configured in the **auto** or **desirable** modes; interfaces configured in the **on** mode do not exchange PAgP packets.

Table 4-21: Modes

EtherChannel Moded	Description
auto	Places an interface into a passive negotiating state, in which the interface responds to PAgP packets it receives but does not start PAgP packet negotiation. This setting minimizes the transmission of PAgP packets.
desirable	Places an interface into an active negotiating state, in which the interface starts negotiations with other interfaces by sending PAgP packets.
on	Forces the interface to channel without PAgP. With the on mode, a usable EtherChannel exists only when an interface group in the on mode is connected to another interface group in the on mode.

Both the **auto** and **desirable** modes allow interfaces to negotiate with partner interfaces to determine if they can form an EtherChannel based on criteria such as interface speed and, for Layer 2 EtherChannels, trunking state and VLAN numbers.

Interfaces can form an EtherChannel when they are in different PAgP modes as long as the modes are compatible. For example:

- An interface in the **desirable** mode can form an EtherChannel with another interface that is in the **desirable** or **auto** mode.
- An interface in the **auto** mode can form an EtherChannel with another interface in the **desirable** mode.

An interface in the **auto** mode cannot form an EtherChannel with another interface that is also in the **auto** mode because neither interface starts PAgP negotiation.

An interface in the **on** mode that is added to a port channel is forced to have the same characteristics as the already existing **on** mode interfaces in the channel.

Note You should exercise care when setting the mode to **on** (manual configuration). All ports configured in the **on** mode are bundled in the same group and are forced to have similar characteristics. If the group is misconfigured, packet loss or spanning-tree loops might occur.

If your switch is connected to a partner that is PAgP-capable, you can configure the switch interface for nonsilent operation by using the **non-silent** keyword. If you do not specify **non-silent** with the **auto** or **desirable** mode, silent mode is assumed.

The silent mode is used when the switch is connected to a device that is not PAgP-capable and seldom, if ever, sends packets. An example of a silent partner is a file server or a packet analyzer that is not generating traffic. In this case, running PAgP on a physical port connected to a silent partner prevents that switch port from ever becoming operational; however, the silent setting allows PAgP to operate, to attach the interface to a channel group, and to use the interface for transmission.

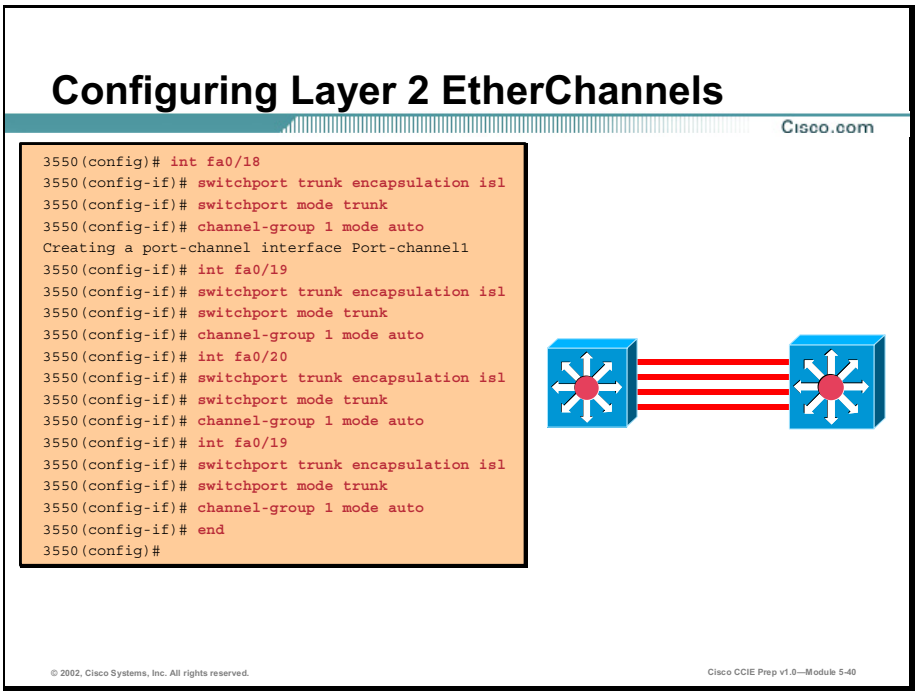

You configure Layer 2 EtherChannels by configuring the Ethernet interfaces with the **channel-group** interface configuration command, which creates the port-channel logical interface.

Use the steps outlined in the following table to assign a Layer 2 Ethernet interface to a Layer 2 EtherChannel:

Table 4-22: Layer 2 EtherChannel

Command	Purpose
3550(config)# **interface** *interface-id*	Enters interface configuration mode and specifies a physical interface to configure. Up to eight interfaces of the same type and speed can be configured for the same group.
3550(config-if)# **switchport mode** {**access** | **trunk**} 3550(config-if)# **switchport access vlan vlan-id**	Assigns the interface as a static-access port in one VLAN or configures it as a trunk. If you configure the interface as a static-access port, assign it to only one VLAN. The range is 1 to 4094.
3550(config-if)# **channel-group** *channel-group-number* **mode** {**auto** [**non-silent**] | **desirable** [**non-silent**] | **on**}	Assigns the interface to a channel group, and specifies the PAgP mode. The default mode is auto silent. For channel-group-number, the range is 1 to 64. Each EtherChannel can have of up to eight compatibly configured Ethernet interfaces.

To remove an interface from the EtherChannel group, use the no channel-group interface configuration command.

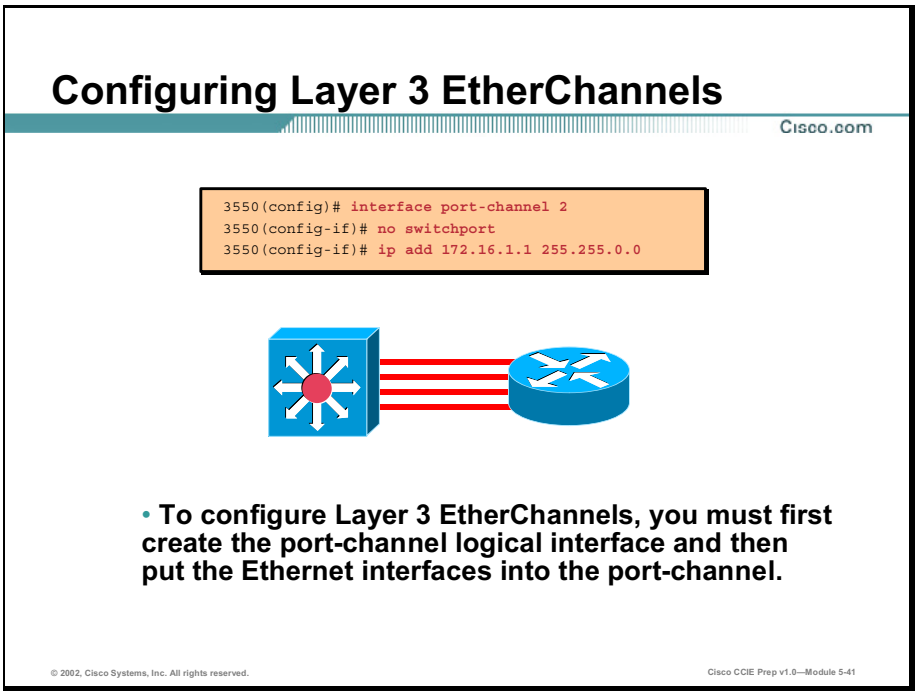

To configure Layer 3 EtherChannels, you must first create the port-channel logical interface and then put the Ethernet interfaces into the port-channel.

Use the steps outlined in the following table to create a port-channel interface for a Layer 3 EtherChannel:

Table 4-23: Layer 3 EtherChannel

Command	Purpose
3550(config)# **interface port-channel** *port-channel-number*	Enters interface configuration mode and creates the port-channel logical interface. For *port-channel-number*, the range is 1 to 64.
3550(config-if)# **no switchport**	Puts the interface into Layer 3 mode.
3550(config-if)# **ip address ip-address mask**	Assigns an IP address and subnet mask to the EtherChannel.

Note To remove the port-channel, use the **no interface port-channel** *port-channel-number* global configuration command.

Note To move an IP address from a physical interface to an EtherChannel, you must delete the IP address from the physical interface before configuring it on the port-channel interface.

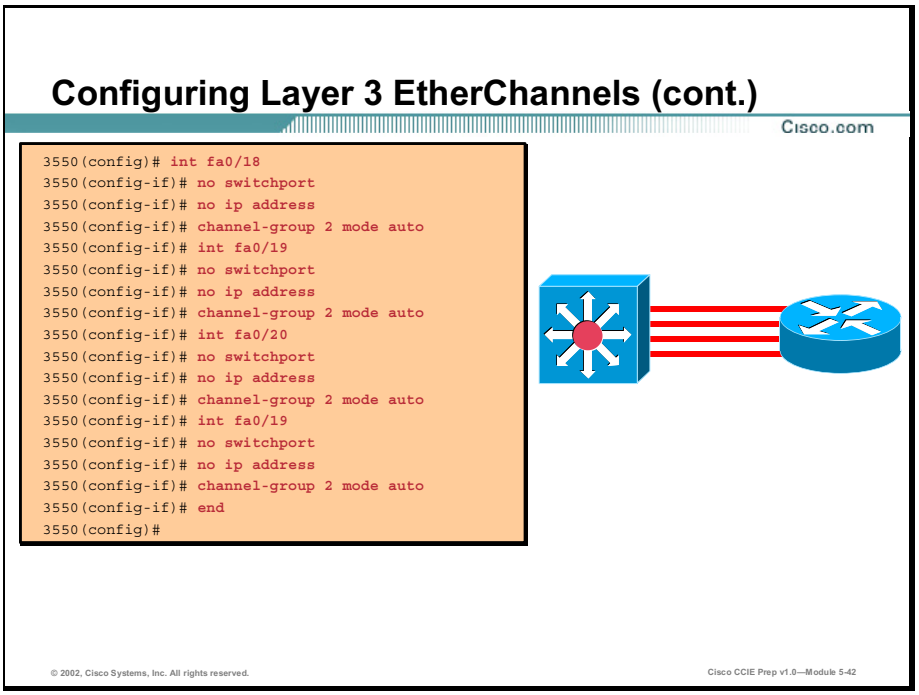

To configure Layer 3 EtherChannels, you must first create the port-channel logical interface and then put the Ethernet interfaces into the port-channel.

Use the steps outlined in the following table to assign an Ethernet interface to a Layer 3 EtherChannel:

Table 4-24: Ethernet Interface

Command	Purpose
3550(config)# **interface** *interface-id*	Enters interface configuration mode and specifies a physical interface to configure. Only physical interfaces can be part of an Etherchannel. Up to eight interfaces of the same type and speed can be configured for the same group.
3550(config-if)# **no ip address**	Ensures that there is no IP address assigned to the physical interface.
3550(config-if)# **channel-group** *channel-***group-number mode** {**auto** [**non-silent**] | **desirable** [**non-silent**] | **on**}	Assigns the interface to a channel group, and specifies the PAgP mode (the default mode is auto silent). For *channel-group-number*, the range is 1 to 64. This number must be the same as the *port-channel-num*ber (logical port) previously configured. Each EtherChannel can consist of up to eight compatibly configured Ethernet interfaces.

Note	To remove an interface from the EtherChannel group, use the **no channel-group** interface configuration command.

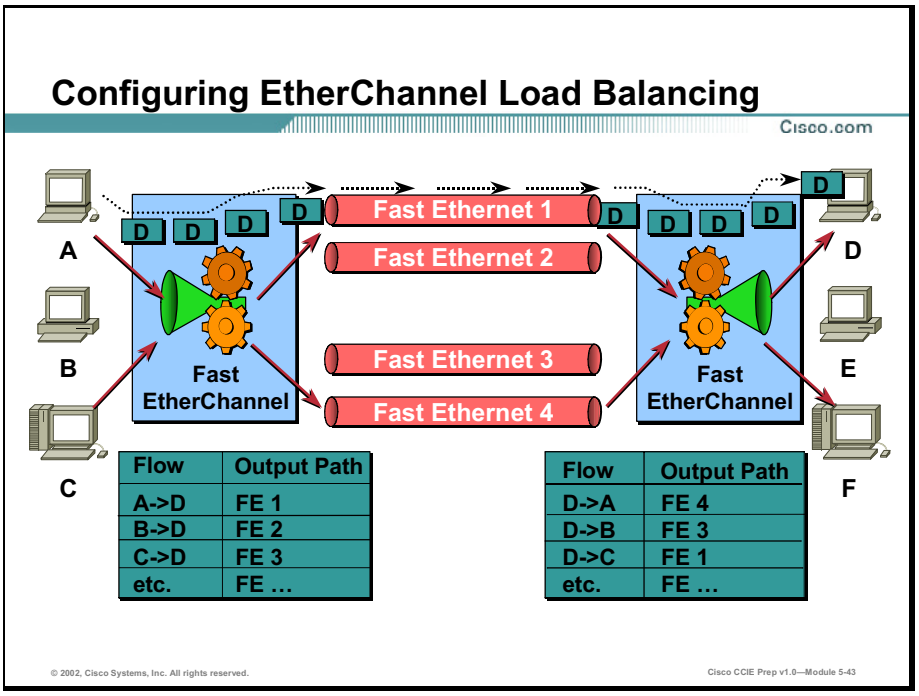

EtherChannel balances the traffic load across the bundled links based on the first two binary bits of a host's MAC address and/or IP address. EtherChannel load balancing can use either source-MAC or destination-MAC address forwarding.

With source-MAC address forwarding, when packets are forwarded to an EtherChannel, they are distributed across the ports in the channel based on the source-MAC address of the incoming packet. Therefore, to provide load balancing, packets from different hosts use different ports in the channel, but packets from the same host use the same port in the channel (and the MAC address learned by the switch does not change).

When source-MAC address forwarding is used, load distribution based on the source and destination IP address is also enabled for routed IP traffic. All routed IP traffic chooses a port based on the source and destination IP address. Packets between two IP hosts always use the same port in the channel, and traffic between any other pair of hosts can use a different port in the channel.

With destination-MAC address forwarding, when packets are forwarded to an EtherChannel, they are distributed across the ports in the channel based on the destination host's MAC address of the incoming packet. Therefore, packets to the same destination are forwarded over the same port, and packets to a different destination are sent on a different port in the channel.

Use the steps outlined in the following table to configure EtherChannel load balancing:

Table 4-25: EtherChannel Load Balancing

Command	Purpose
3550(config)# **port-channel load-balance** {**dst-mac** | **src-mac**}	Configures an EtherChannel load-balancing method: ■ **dst-mac**—Load distribution is based on the destination-host MAC address of the incoming packet. Packets to the same destination are sent on the same port, but packets to different destinations are sent on different ports in the channel. src-mac—Load distribution is based on the source-MAC address of the incoming packet. Packets from different hosts use different ports in the channel, but packets from the same host use the same port in the channel. The default is src-mac

Note To return EtherChannel load balancing to the default configuration, use the **no port-channel load-balance** global configuration command.

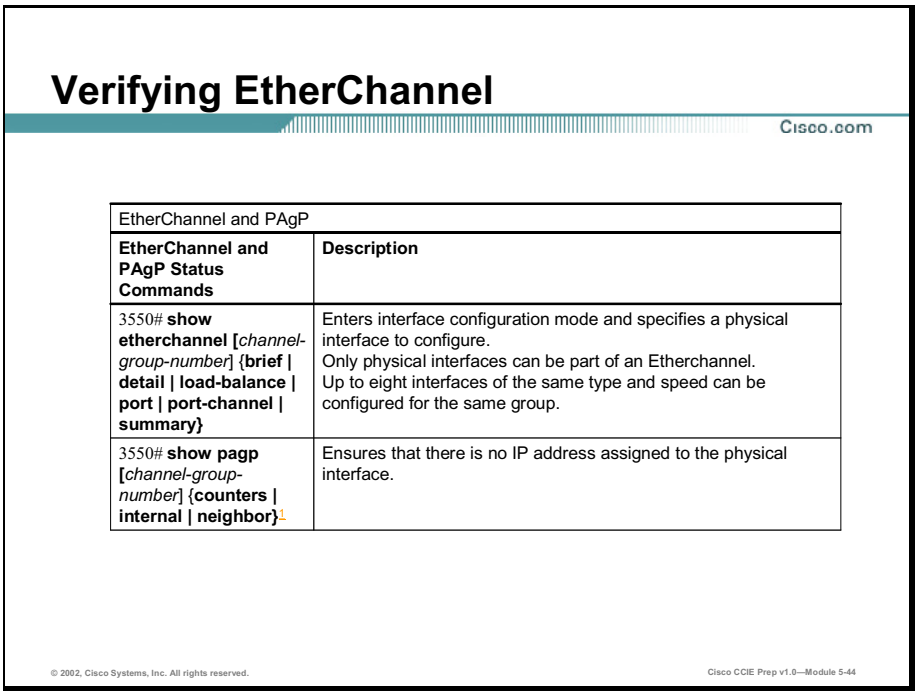

Thefollowing table lists the commands used to display EtherChannel and PAgP status information:

Table 4-26: EtherChannel and PAgP

EtherChannel and PAgP Status Commands	Description
3550# **show etherchannel** [*channel-group-number*] {**brief	detail	load-balance	port	port-channel	summary**}	Enters interface configuration mode and specifies a physical interface to configure. Only physical interfaces can be part of an Etherchannel. Up to eight interfaces of the same type and speed can be configured for the same group.
3550# **show pagp** [*channel-group-number*] {**counters	internal	neighbor**}[1]	Ensures that there is no IP address assigned to the physical interface.

Summary

This section summarizes the key points discussed in this lesson.

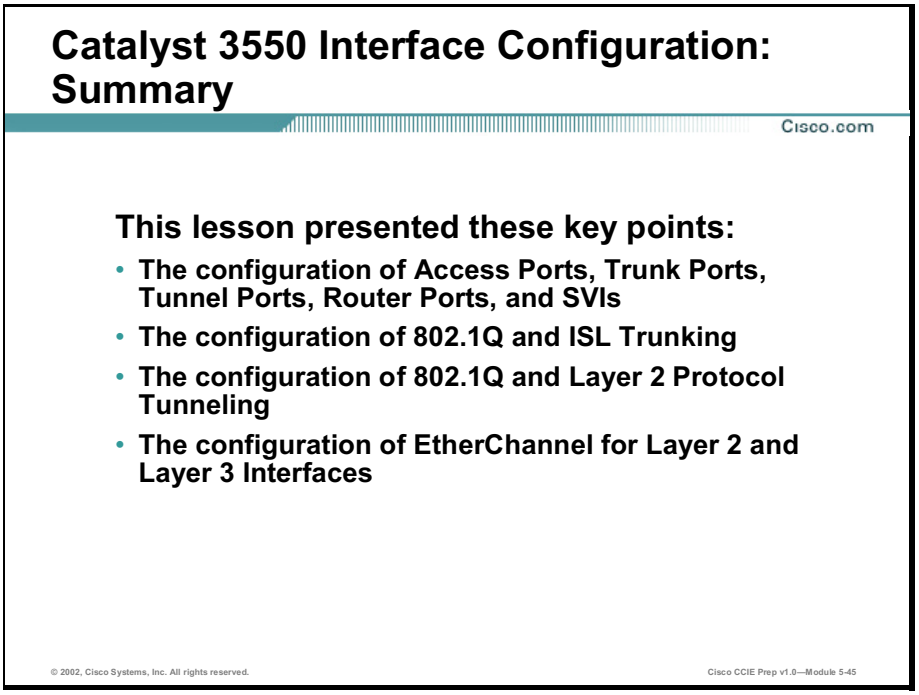

Next Steps

After completing this lesson, go to:

- Catalyst 3550 Advanced Configuration

References

For additional information, refer to these resources:

- *CCIE Professional Development : Cisco Lan Switching* by Kennedy Clark

Lesson Assessment (Quiz)

Q1) Which of the following are valid switch port types on the Catalyst 3550?

A) Trunk Ports

B) Tunnel Ports

C) VLAN Ports

D) Hybrid Ports

E) Access Ports

Q2) List the two commands that are required in interface configuration mode to make a switch port an access port.

switchport mode access

switchport access vlan <vlan id>

Q3) Which of the following commands is used to specify the native vlan on an 802.1Q trunk?

A) **switchport dot1q native** <vlan id>

B) **switchport dot1q trunk native** <vlan id>

C) **dot1q trunk native** <vlan id>

D) **switchport trunk native vlan** <vlan id>

Q4) The Catalyst 3550 supports which of the following tunneling mechanisms?

A) PPTP

B) IPSec

C) 802.1Q Tunneling

D) Layer 2 Protocol Tunneling

Q5) List the command used in interface configuration mode to turn a Layer 2 switch port into a Layer 3 router port.

no switchport

Which of the following protocols facilitates the automatic creation of EtherChannels?

A) Dynamic Trunk Protocol (DTP)

B) VLAN Trunking Protocol (VTP)

C) Port Aggregation Protocol (PAgP)

D) None of the above (EtherChannels must be manually created)

Catalyst 3550 Advanced Configuration

Overview

The Catalyst 3550 is an advanced next-generation Multilayer switch. It can perform Layer 2 and Layer 3 functions. It also supports advanced security and QoS features.

Importance

In addition to configuring the Catalyst 3550 for basic operation, you will also be asked to configure some advanced features on the Catalyst 3550.

Objectives

Upon completing this lesson, you will be able to configure:

- Understand the concepts of Spanning Tree
- Fine tune Spanning Tree to optimize convergence after a network failure
- Monitor and analyze network traffic using SPAN and RSPAN
- Configure Fallback Bridging to bridge non-IP protocols

Learner Skills and Knowledge

To fully benefit from this lesson, you must have these prerequisite skills and knowledge:

- Passed the CCIE written qualification exam
- Completed the Building Cisco Multilayer Switched Networks (BCMSN) course, or have the equivalent knowledge

Outline

This lesson includes these sections:

- Overview
- Spanning Tree Operation
- Monitoring and Analyzing Traffic
- Fallback Bridging
- Summary
- Lesson Assessment (Quiz)

Spanning Tree

Spanning tree defaults may be modified to improve performance on your network. This section will focus on tuning the parameters of Spanning Tree on the Catalyst 3550 switch.

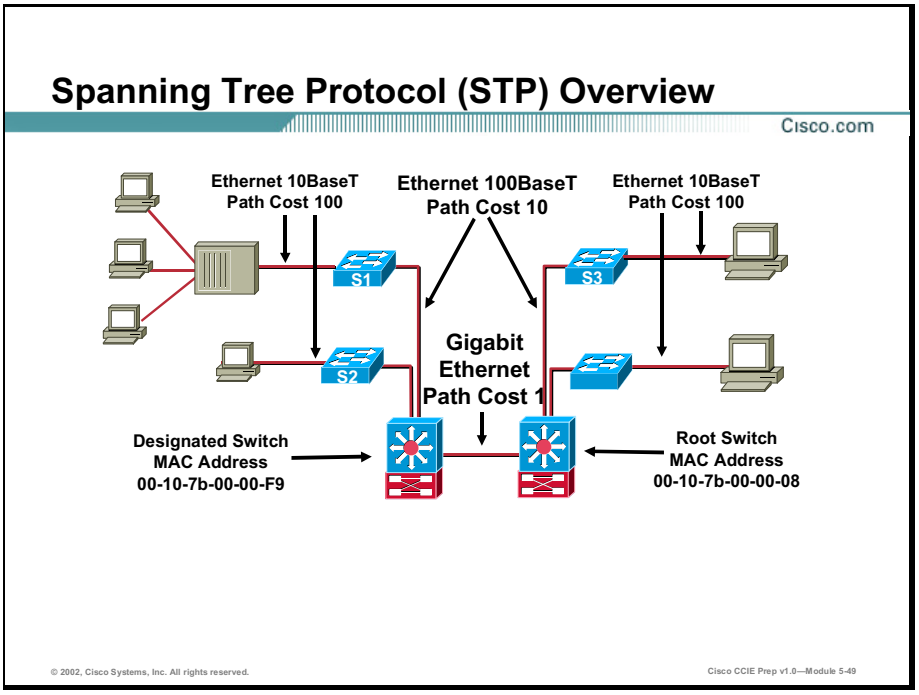

Spanning Tree Protocol (STP) is used as a distributed algorithm to create one path in a redundant layer 2 network per VLAN. The algorithm relies on a root bridge (switch) per VLAN to aid in generating one path per VLAN and transmissions of Bridge Protocol Data Units (BPDUs). Each port on a switch using STP exists in one of the following five states:

- Blocking: A port in the blocking state does not participate in frame forwarding
- Listening: The listening state is the first transitional state a port enters after the blocking state; in this state, a port only listens for BPDUs
- Learning: In the learning state, the port prepares for frame forwarding by adding MAC addresses to the CAM
- Forwarding: In the forwarding state, the port forwards frames
- Disabled: In the disabled state, a port is virtually non-operational

A port moves through these five states as follows:

- From initialization to blocking
- From blocking to listening or to disabled

- From listening to learning or to disabled
- From learning to forwarding or to disabled
- From forwarding to disabled in the event a loop is detected

There are default timers associated with each state. STP uses these timers to determine one path through a redundant network.

Table 5-27: Timers

Timer	Description	Default (Sec)
Hello time	Determines how often the switch broadcasts Hello messages to other switches	2
Forward delay time	Determines the amount of time a port will remain in the listening and learning states before entering the forwarding state	15
Maximum age time	Determines the amount of time protocol information received on a port is stored by the switch	20

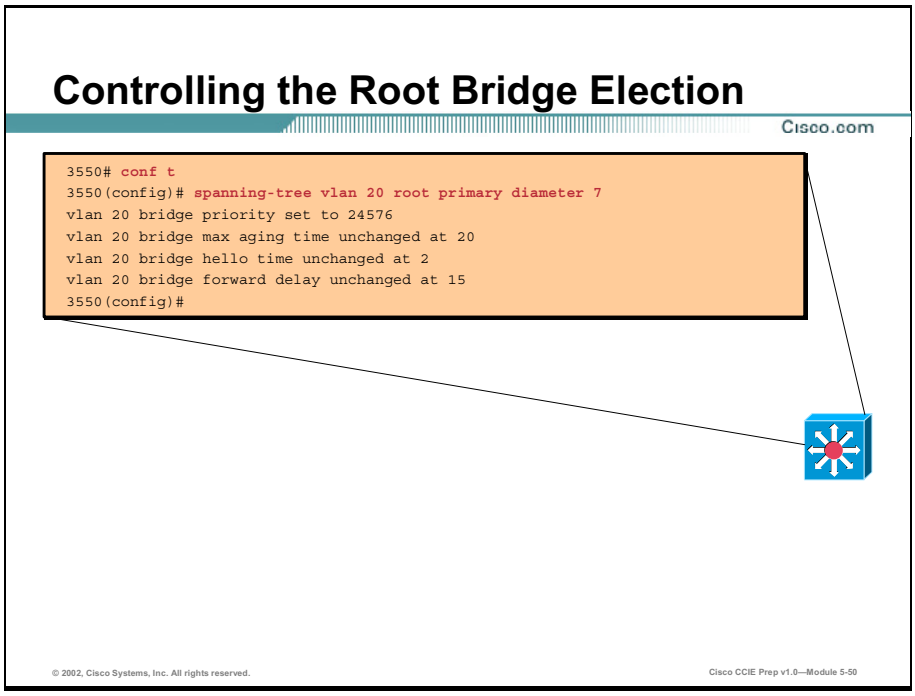

The switch maintains a separate spanning-tree instance for each active VLAN configured on it. A bridge ID, consisting of the switch priority and the switch MAC address, is associated with each instance. For each VLAN, the switch with the lowest bridge ID becomes the root switch for that VLAN.

To configure a switch to become the root for the specified VLAN, use the **spanning-tree vlan** *vlan-id* **root** global configuration command to modify the switch priority from the default value (32768) to a significantly lower value. When you enter this command, the switch checks the switch priority of the root switches for each VLAN. Because of the extended system ID support, the switch sets its own priority for the specified VLAN to 24576 if this value will cause this switch to become the root for the specified VLAN.

If any root switch for the specified VLAN has a switch priority lower than 24576, the switch sets its own priority for the specified VLAN to 4096 less than the lowest switch priority. 4096 is the value of the least-significant bit of a 4-bit switch priority value.

Note	The **spanning-tree vlan** *vlan-id* **root** global configuration command fails if the value necessary to be the root switch is less than 1.

Use the steps outlined in the following table to configure the switch to become the root for the specified VLAN:

Table 5-28: Root for VLAN

Command	Purpose
spanning-tree vlan *vlan-id* **root primary** [**diameter** *net-diameter*]	Configure a switch to become the root for the specified VLAN. ■ For *vlan-id*, the range is 1 to 4094. Do not enter leading zeros. (Optional) For **diameter** *net-diameter*, specify the maximum number of switches between any two end stations. The range is 2 to 7.

To return the switch to its default setting, use the **no spanning-tree vlan** *vlan-id* **root** global configuration command.

Use the **diameter** keyword to specify the Layer 2 network diameter (that is, the maximum number of switch hops between any two end stations in the Layer 2 network). When you specify the network diameter, the switch automatically sets an optimal hello time, forward-delay time, and maximum-age time for a network of that diameter, which can significantly reduce the convergence time.

Note After configuring the switch as the root switch, we recommend that you avoid manually configuring the hello time, forward-delay time, and maximum-age time by using the **spanning-tree** vlan *vlan-id* **hello-time**, **spanning-tree** vlan *vlan-id* **forward-time**, and the **spanning-tree** vlan *vlan-id* **max-age** global configuration commands. Use the **diameter** keyword instead.

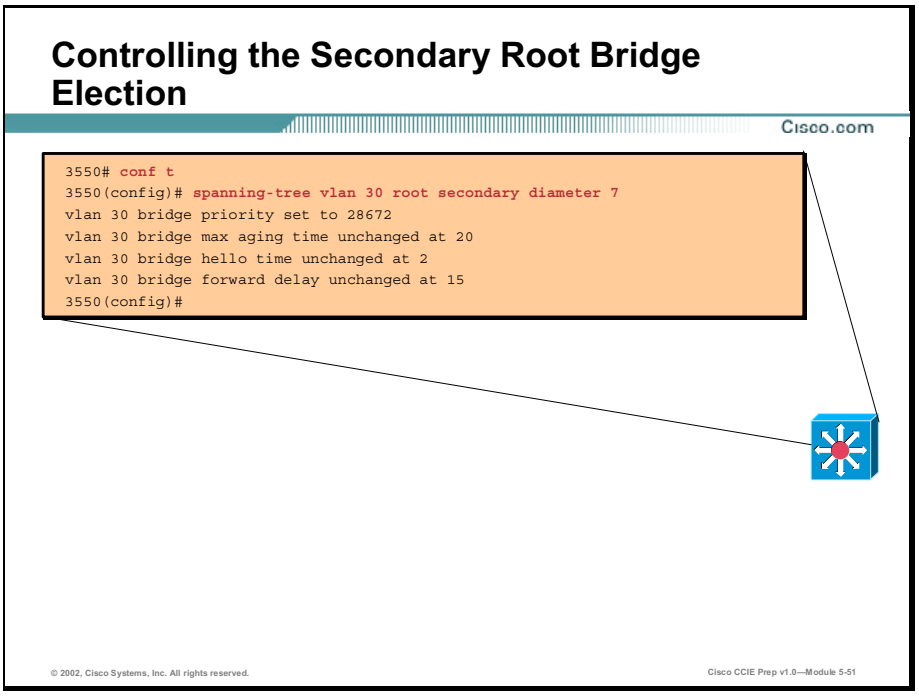

You can also configure a Catalyst 3550 switch that supports the extended system ID as the secondary root bridge. This action modifies the switch's bridge priority from the default value (32768) to 28672. The switch is then likely to become the root switch for the specified VLAN if the primary root switch fails. This is assuming that the other network switches use the default switch priority of 32768 and therefore are unlikely to become the root switch. For Catalyst 3550 switches without the extended system ID support (software earlier than Release 12.1(8)EA1), the switch priority is changed to 16384.

You can execute this command on more than one switch to configure multiple backup root switches. Use the same network diameter as you used when you configured the primary root switch.

Use the steps outline in the following table to configure the switch to become the secondary root for the specified VLAN:

Table 5-29: Root for VLAN

Command	Purpose
spanning-tree vlan *vlan-id* **root secondary [diameter** *net-diameter*]	Configure a switch to become the secondary root for the specified VLAN. ■ For *vlan-id*, the range is 1 to 4094. Do not enter leading zeros. ■ (Optional) For **diameter** *net-diameter*, specify the maximum number of switches between any two end stations. The range is 2 to 7. Use the same network diameter that you used when configuring the primary root switch.

To return the switch to its default setting, use the **no spanning-tree vlan** *vlan-id* **root** global configuration command.

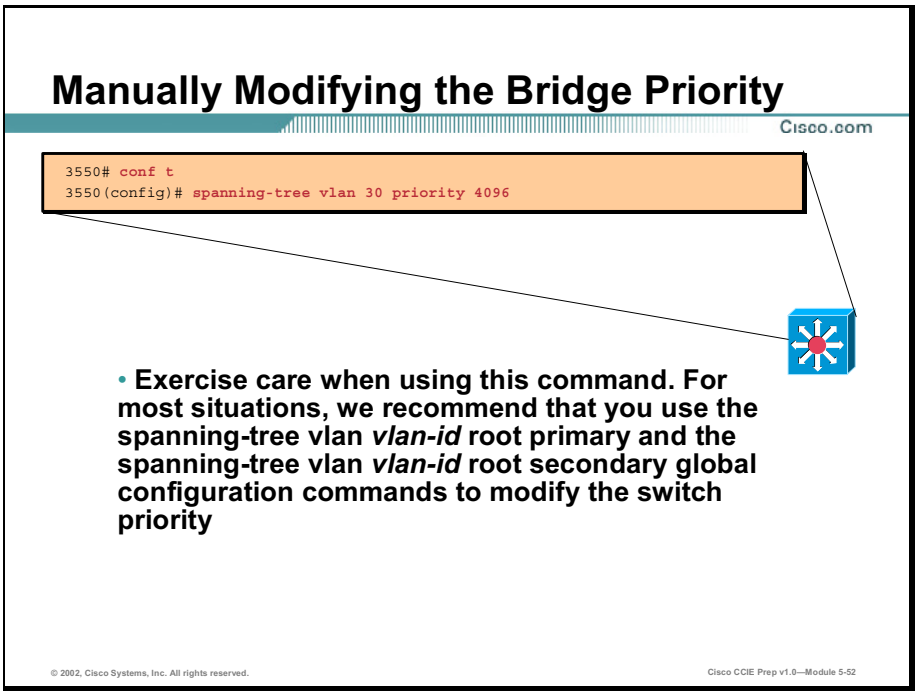

You can also manually configure the switch's bridge priority and make it more likely that the switch will be chosen as the root switch for a particular VLAN. Manually modifying the bridge priority of a switch can have undesired affects on your network, therefore it is recommended to use the **spanning-tree vlan** *vlan-id* **root primary** and the **spanning-tree vlan** *vlan-id* **root secondary** commands to systematically control the root bridge election.

Use the steps outlined in the table below to manually configure the switch's bridge priority for a particular VLAN:

Table 5-30: Root for VLAN

Command	Purpose
spanning-tree vlan *vlan-id* **priority** *priority*	Configure the switch priority of a VLAN. ■ For *vlan-id*, the range is 1 to 4094. Do not enter leading zeros. ■ For *priority*, the range is 0 to 61440 in increments of 4096; the default is 32768. The lower the number, the more likely the switch will be chosen as the root switch. Valid priority values are 4096, 8192, 12288, 16384, 20480, 24576, 28672, 32768, 36864, 40960, 45056, 49152, 53248, 57344, and 61440. All other values are rejected.

To return the switch to its default setting, use the no spanning-tree vlan *vlan-id* priority global configuration command.

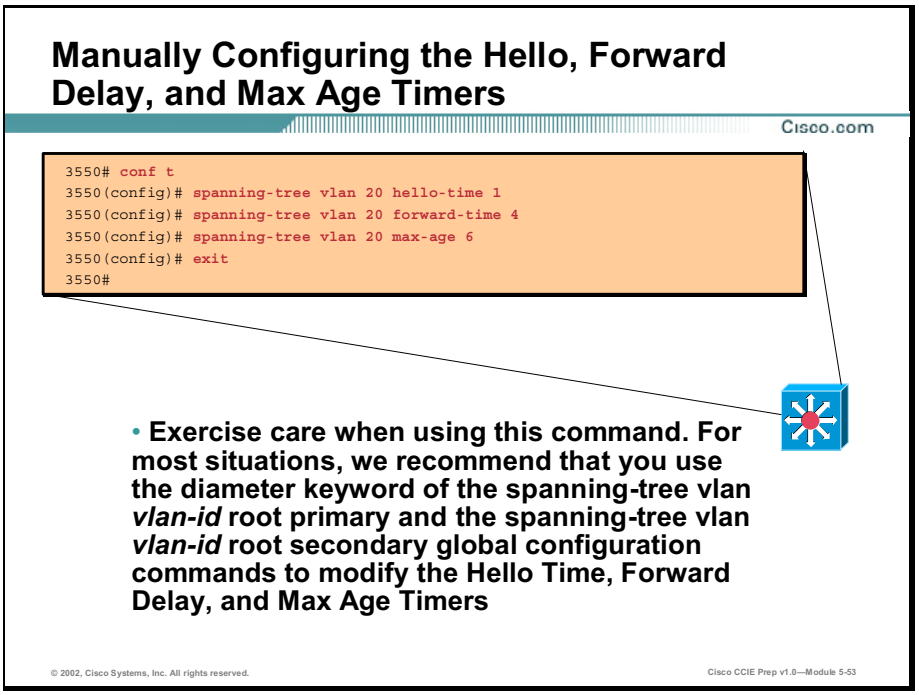

Although it is recommended to use the diameter keyword of **spanning-tree vlan** *vlan-id* **root** command to control this values, they can be manually set. Use the following commands to manually configure the Hello, Forward Delay, and Max Age Timers.

Table 5-31: Configure Hello, Foreward Delay, and Max Age Timers

Command	Purpose
spanning-tree vlan *vlan-id* **hello-time** *seconds*	Configure the hello time of a VLAN. The hello time is the interval between the generation of configuration messages by the root switch. These messages mean that the switch is alive. ■ For *vlan-id*, the range is 1 to 4094. Do not enter leading zeros. ■ For *seconds*, the range is 1 to 10; the default is 2.
spanning-tree vlan *vlan-id* **forward-time** *seconds*	Configure the forward time of a VLAN. The forward delay is the number of seconds a port waits before changing from its spanning-tree learning and listening states to the forwarding state. ■ For *vlan-id*, the range is 1 to 4094. Do not enter leading zeros. ■ For *seconds*, the range is 4 to 30; the default is 15.
spanning-tree vlan *vlan-id* **max-age** *seconds*	Configure the maximum-aging time of a VLAN. The maximum-aging time is the number of seconds a switch waits without receiving spanning-tree configuration messages before attempting a reconfiguration. ■ For *vlan-id*, the range is 1 to 4094. Do not enter leading zeros. ■ For *seconds*, the range is 6 to 40; the default is 20.

To return the switch to its default settings, use the **no spanning-tree vlan** *vlan-id* **hello-time**, **no spanning-tree vlan** *vlan-id* **forward-time**, and **no spanning-tree vlan** *vlan-id* **max-age** commands respectively.

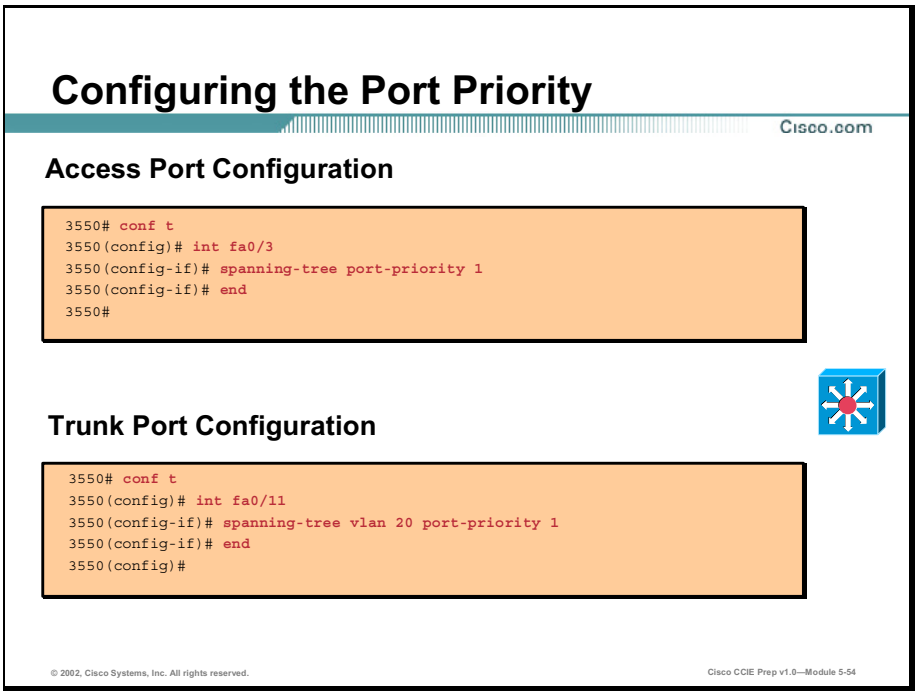

If a loop occurs, spanning tree uses the port priority when selecting an interface to put into the forwarding state. You can assign higher priority values (lower numerical values) to interfaces that you want selected first and lower priority values (higher numerical values) that you want selected last. If all interfaces have the same priority value, spanning tree puts the interface with the lowest interface number in the forwarding state and blocks the other interfaces.

Depending on the type of port you want to configure, they are two different ways to control the port priority. The Catalyst 3550 uses the port priority value when the interface is configured as an access port and uses VLAN port priority values when the interface is configured as a trunk port.

Use the steps outlined in the following table to configure the port priority of an interface:

Table 5-32: Port Priority

Command	Purpose
interface *interface-id*	Enter interface configuration mode, and specify an interface to configure. Valid interfaces include physical interfaces and port-channel logical interfaces (**port-channel** *port-channel-number*).
spanning-tree port-priority *priority*	Configure the port priority for an interface that is an access port. For *priority*, the range is 0 to 255; the default is 128. The lower the number, the higher the priority.
spanning-tree vlan *vlan-id* **port-priority** *priority*	Configure the VLAN port priority for an interface that is a trunk port. ■ For *vlan-id*, the range is 1 to 4094. Do not enter leading zeros. ■ For *priority*, the range is 0 to 255; the default is 128. The lower the number, the higher the priority.

To return the interface to its default setting, use the **no spanning-tree** [**vlan** *vlan-id*] **port-priority** interface configuration command.

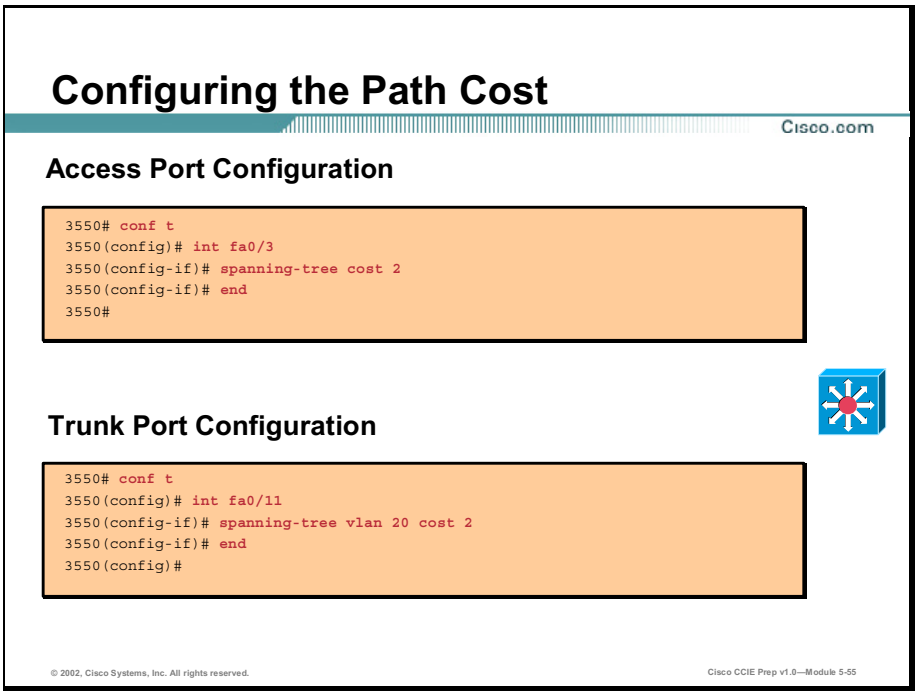

The spanning-tree path cost default value is derived from the media speed of an interface. If a loop occurs, spanning tree uses cost when selecting an interface to put in the forwarding state. You can assign lower cost values to interfaces that you want selected first and higher cost values that you want selected last. If all interfaces have the same cost value, spanning tree puts the interface with the lowest interface number in the forwarding state and blocks the other interfaces.

Spanning tree uses the cost value when the interface is configured as an access port and uses VLAN port cost values when the interface is configured as a trunk port.

Use the steps outlined in the following table to configure the cost of an interface:

Table 5-33: Cost of an Interface

Command	Purpose
interface *interface-id*	Enter interface configuration mode, and specify an interface to configure. Valid interfaces include physical interfaces and port-channel logical interfaces (**port-channel** *port-channel-number*).
spanning-tree cost *cost*	Configure the cost for an interface that is an access port. If a loop occurs, spanning tree uses the path cost when selecting an interface to place into the forwarding state. A lower path cost represents higher-speed transmission. For *cost*, the range is 1 to 200000000; the default value is derived from the media speed of the interface.
spanning-tree vlan *vlan-id* **cost** *cost*	Configure the VLAN cost for an interface that is a trunk port. If a loop occurs, spanning tree uses the path cost when selecting an interface to place into the forwarding state. A lower path cost represents higher-speed transmission. ■ For *vlan-id*, the range is 1 to 4094. Do not enter leading zeros. ■ For *cost*, the range is 1 to 200000000; the default value is derived from the media speed of the interface.

To return the interface to its default setting, use the **no spanning-tree** [**vlan** *vlan-id*] **cost** interface configuration command.

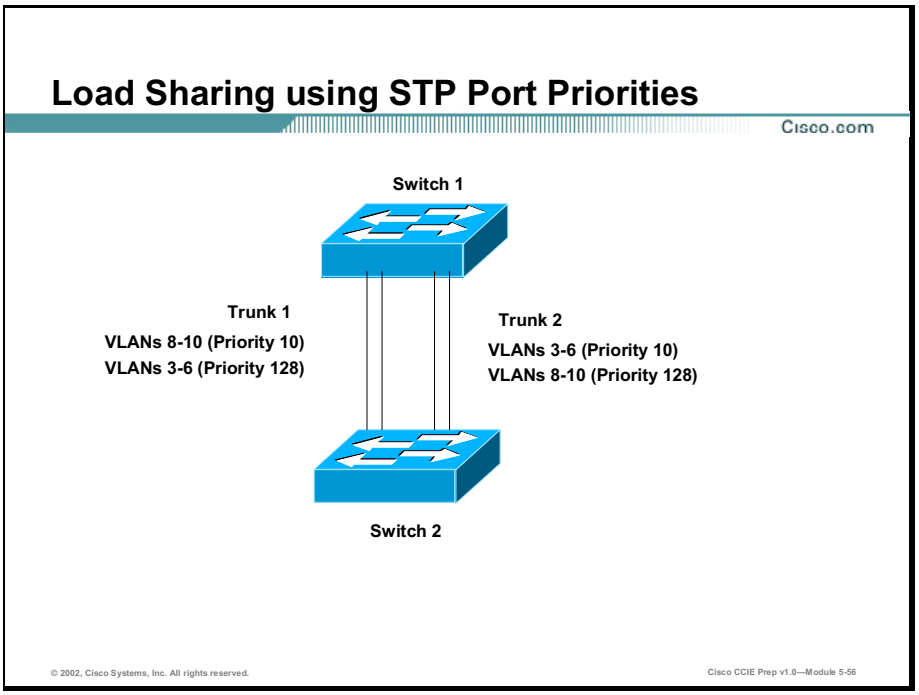

Load sharing divides the bandwidth supplied by parallel trunks connecting switches. To avoid loops, STP normally blocks all but one parallel link between switches. Using load sharing, you divide the traffic between the links according to which VLAN the traffic belongs.

You configure load sharing on trunk ports by using STP port priorities or STP path costs. For load sharing using STP port priorities, both load-sharing links must be connected to the same switch. For load sharing using STP path costs, each load-sharing link can be connected to the same switch or to two different switches.

Load Sharing Using STP Port Priorities

When two ports on the same switch form a loop, the STP port priority setting determines which port is enabled and which port is in a blocking state. You can set the priorities on a parallel trunk port so that the port carries all the traffic for a given VLAN. The trunk port with the higher priority (lower values) for a VLAN is forwarding traffic for that VLAN. The trunk port with the lower priority (higher values) for the same VLAN remains in a blocking state for that VLAN. One trunk port sends or receives all traffic for the VLAN.

The figure above shows two trunks connecting supported switches. In this example, the switches are configured as follows:

- VLANs 8 through 10 are assigned a port priority of 10 on Trunk 1.
- VLANs 3 through 6 retain the default port priority of 128 on Trunk 1.
- VLANs 3 through 6 are assigned a port priority of 10 on Trunk 2.
- VLANs 8 through 10 retain the default port priority of 128 on Trunk 2.

In this example, Trunk 1 carries traffic for VLANs 8 through 10, and Trunk 2 carries traffic for VLANs 3 through 6. If the active trunk fails, the trunk with the lower priority takes over and carries the traffic for all of the VLANs. No duplication of traffic occurs over any trunk port.

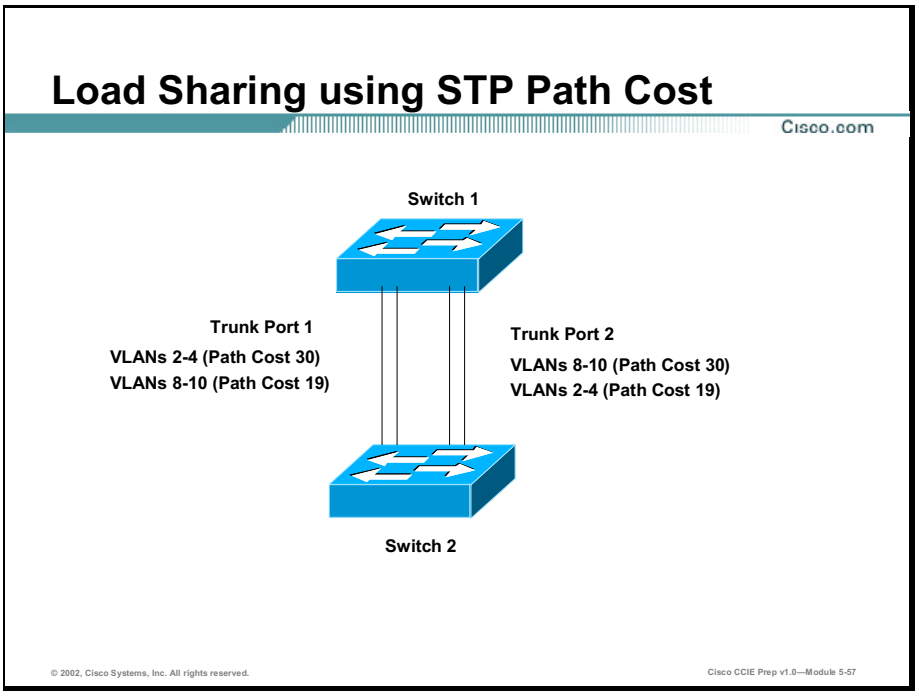

You can configure parallel trunks to share VLAN traffic by setting different path costs on a trunk and associating the path costs with different sets of VLANs. The VLANs keep the traffic separate. Because no loops exist, STP does not disable the ports, and redundancy is maintained in the event of a lost link.

In the figure above, Trunk ports 1 and 2 are 100BASE-T ports. The path costs for the VLANs are assigned as follows:

- VLANs 2 through 4 are assigned a path cost of 30 on Trunk port 1.
- VLANs 8 through 10 retain the default 100BASE-T path cost on Trunk port 1 of 19.
- VLANs 8 through 10 are assigned a path cost of 30 on Trunk port 2.
- VLANs 2 through 4 retain the default 100BASE-T path cost on Trunk port 2 of 19.

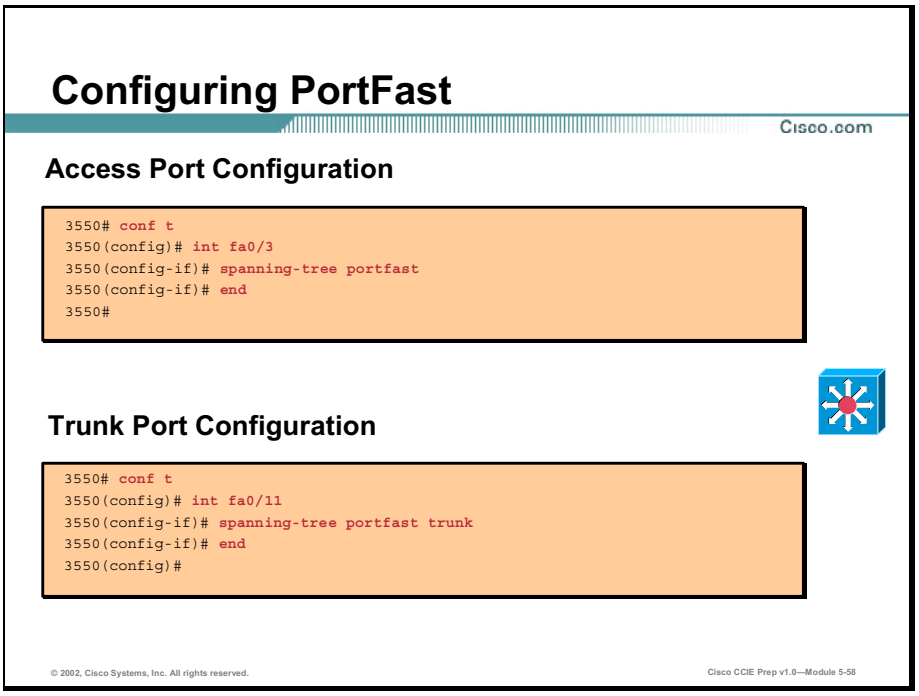

Port Fast immediately brings an interface configured as an access or trunk port to the forwarding state from a blocking state, bypassing the listening and learning states. You can use Port Fast on ports connected to a single workstation or server, to allow those devices to immediately connect to the network, rather than waiting for the spanning tree to converge.

Ports connected to a single workstation or server should not receive bridge protocol data units (BPDUs). A port with Port Fast enabled goes through the normal cycle of spanning-tree status changes when the switch is restarted.

Note	Because the purpose of Port Fast is to minimize the time ports must wait for spanning-tree to converge, it is effective only when used on ports connected to end stations. If you enable Port Fast on a port connecting to another switch, you risk creating a spanning-tree loop.

A port with the Port Fast feature enabled is moved directly to the spanning-tree forwarding state without waiting for the standard forward-time delay.

Caution	Use Port Fast *only* when connecting a single end station to an access or trunk port. Enabling this feature on a port connected to a switch or hub could prevent spanning tree from detecting and disabling loops in your network, which could cause broadcast storms and address-learning problems.

You can enable this feature if your switch is running PVST or MSTP.

Use the steps outlined in the following table to enable PortFast:

Table 5-34: PortFast

Command	Purpose
interface *interface-id*	Enter interface configuration mode, and specify an interface to configure.
spanning-tree portfast [**trunk**]	Enable Port Fast on an access port connected to a single workstation or server. By specifying the **trunk** keyword, you can enable Port Fast on a trunk port. **Caution** Make sure that there are no loops in the network between the trunk port and the workstation or server before you enable Port Fast on a trunk port. By default, Port Fast is disabled on all ports.

Note You can use the **spanning-tree portfast default** global configuration command to globally enable the Port Fast feature on all nontrunking ports.

To disable the Port Fast feature, use the **spanning-tree portfast disable** interface configuration command.

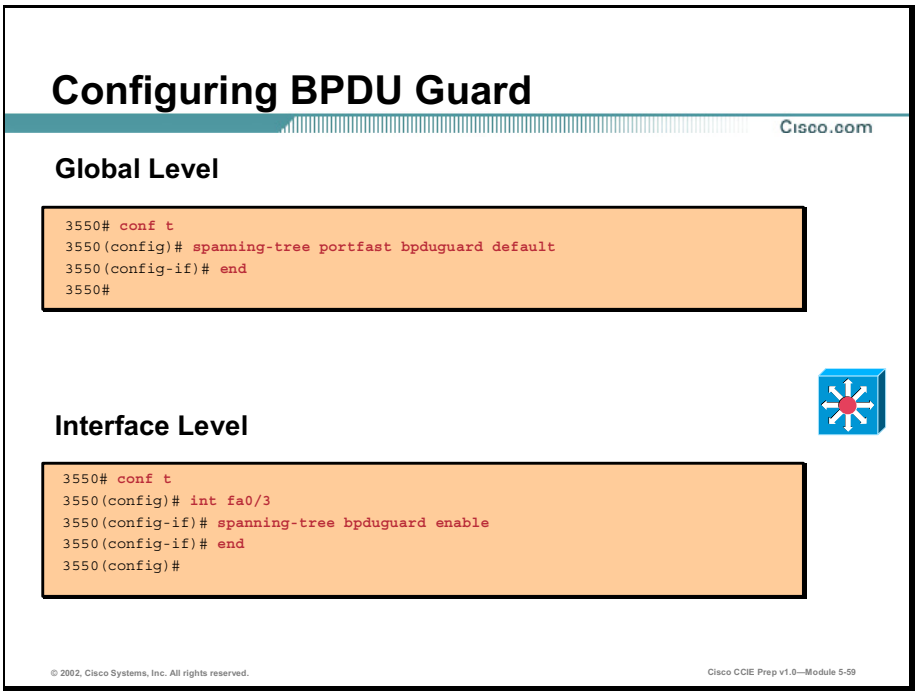

In a valid configuration, Port Fast-enabled ports should not receive BPDUs. Receiving a BPDU on a Port Fast-enabled port signals that the port is connected to a switch and not an end station. This could indicate the connection of an unauthorized switch. Since connecting switches to port fast enabled ports can create a loop in the topology and cause network disruptions, it is critical to have a way to prevent this. The BPDU guard feature is used to monitor the reception of BPDUs on port fast enabled ports.

If BPDU guard is enabled and a BPDU is received on a port fast enabled port, the BPDU guard feature puts the port in the error-disabled state. The BPDU guard feature provides a secure response to invalid configurations because the administrator must manually put the port back in service. The BPDU guard feature can be globally enabled on the switch or can be enabled per interface, but the feature operates with some differences.

At the global level, you can enable BPDU guard on Port Fast-enabled ports by using the **spanning-tree portfast bpduguard default** global configuration command. Spanning tree shuts down ports that are in a Port Fast-operational state. In a valid configuration, Port Fast-enabled ports do not receive BPDUs. Receiving a BPDU on a Port Fast-enabled port signals an invalid configuration, such as the connection of an unauthorized device, and the BPDU guard feature puts the port in the error-disabled state.

At the interface level, you can enable BPDU guard on any port by using the **spanning-tree bpduguard enable** interface configuration command without also enabling the Port Fast feature. When the port receives a BPDU, it is put in the error-disabled state.

When you globally enable BPDU guard on ports that are Port Fast-enabled (the ports are in a Port Fast-operational state), spanning tree shuts down Port Fast-enabled ports that receive BPDUs.

You can also use the **spanning-tree bpduguard enable** interface configuration command to enable BPDU guard on any port without also enabling the Port Fast feature. When the port receives a BPDU, it is put in the error-disabled state.

Use the steps outlined in the following table to enable the BPDU guard feature:

Table 5-35: BPDU Guard

Command	Purpose
spanning-tree portfast bpduguard default	Globally enable BPDU guard. By default, BPDU guard is disabled.
interface *interface-id*	Enter interface configuration mode, and specify the interface connected to an end station.
spanning-tree portfast	Enable the Port Fast feature.

To disable BPDU guard, use the **no spanning-tree portfast bpduguard default** global configuration command.

You can override the setting of the **no spanning-tree portfast bpduguard default** global configuration command by using the **spanning-tree bpduguard enable** interface configuration command.

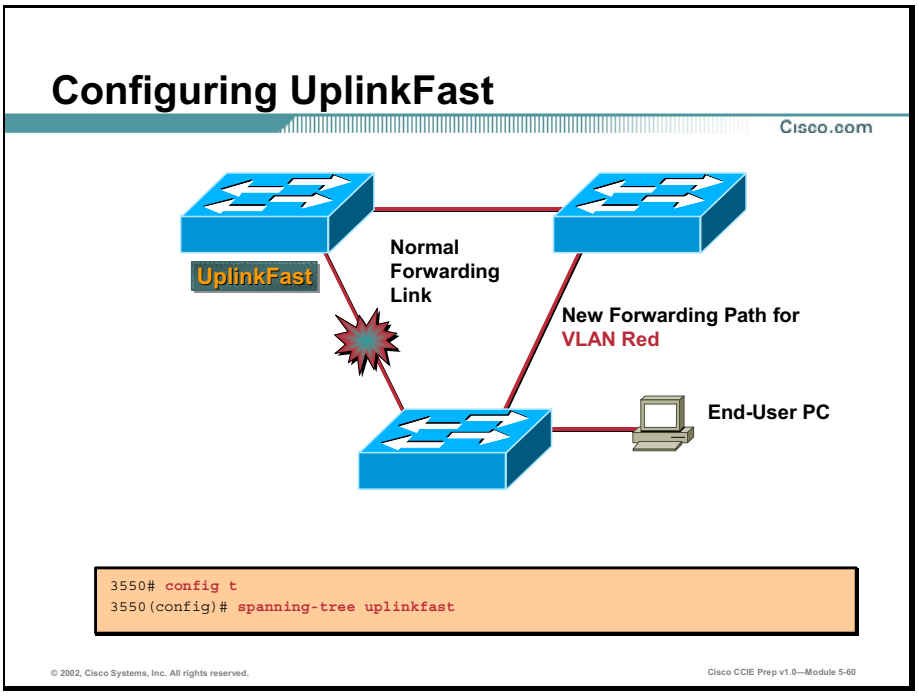

If a switch looses connectivity, it begins using the alternate paths as soon as the spanning tree selects a new root port. By enabling UplinkFast, you can accelerate the choice of a new root port when a link or switch fails or when the spanning tree reconfigures itself. The root port transitions to the forwarding state immediately without going through the listening and learning states, as it would with the normal spanning-tree procedures. The UplinkFast feature is supported only when the switch is running PVST.

Note	UplinkFast is most useful in wiring-closet switches at the access or edge of the network. It is not appropriate for backbone devices. This feature might not be useful for other types of applications.

UplinkFast provides fast convergence after a direct link failure and achieves load balancing between redundant Layer 2 links using uplink groups. An uplink group is a set of Layer 2 interfaces (per VLAN), only one of which is forwarding at any given time. Specifically, an uplink group consists of the root port (which is forwarding) and a set of blocked ports, except for self-looping ports. The uplink group provides an alternate path in case the currently forwarding link fails.

UplinkFast cannot be enabled on VLANs that have been configured for switch priority. To enable UplinkFast on a VLAN with switch priority configured, first restore the switch priority on the VLAN to the default value by using the **no spanning-tree vlan** *vlan-id* **priority** global configuration command.

Note	When you enable UplinkFast, it affects all VLANs on the switch. You cannot configure UplinkFast on an individual VLAN.

The UplinkFast feature is supported only when the switch is running PVST.

Use the steps outlined in the following table to enable UplinkFast:

Table 5-36: UplinkFast

Command	Purpose
spanning-tree uplinkfast [**max-update-rate** *pkts-per-second*]	Enable UplinkFast. (Optional) For *pkts-per-second*, the range is 0 to 65535 packets per second; the default is 150. If you set the rate to 0, station-learning frames are not generated, and the spanning-tree topology converges more slowly after a loss of connectivity.

When UplinkFast is enabled, the switch priority of all VLANs is set to 49152. If you change the path cost to a value less than 3000 and you enable UplinkFast or UplinkFast is already enabled, the path cost of all interfaces and VLAN trunks is increased by 3000 (if you change the path cost to 3000 or above, the path cost is not altered). The changes to the switch priority and the path cost reduces the chance that the switch will become the root switch.

To return the update packet rate to the default setting, use the **no spanning-tree uplinkfast max-update-rate** global configuration command. To disable UplinkFast, use the **no spanning-tree uplinkfast** command. When UplinkFast is disabled, the switch priorities of all VLANs and path costs of all interfaces are set to default values if you did not modify them from their defaults.

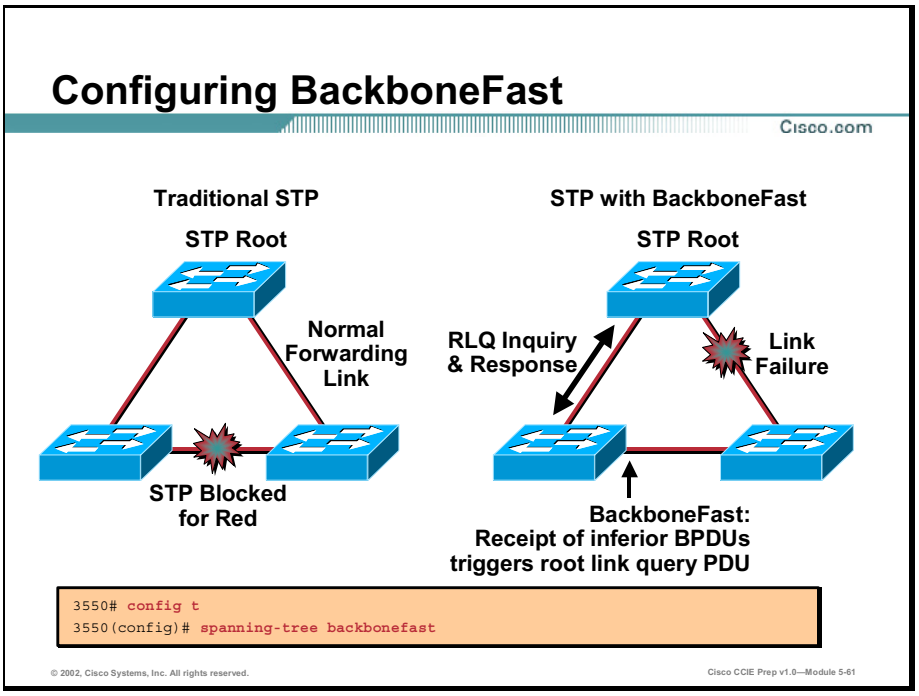

BackboneFast detects indirect failures in the core of the backbone. BackboneFast is a complementary technology to the UplinkFast feature, which responds to failures on links directly connected to access switches. BackboneFast optimizes the maximum-age timer, which determines the amount of time the switch stores protocol information received on an interface. When a switch receives an inferior BPDU from the designated port of another switch, the BPDU is a signal that the other switch might have lost its path to the root, and BackboneFast tries to find an alternate path to the root.

BackboneFast, which is enabled by using the **spanning-tree backbonefast** global configuration command, starts when a root port or blocked port on a switch receives inferior BPDUs from its designated bridge. An inferior BPDU identifies one switch as both the root bridge and the designated bridge. When a switch receives an inferior BPDU, it means that a link to which the switch is not directly connected (an indirect link) has failed (that is, the designated bridge has lost its connection to the root switch). Under spanning-tree rules, the switch ignores inferior BPDUs for the configured maximum aging time specified by the **spanning-tree max-age** global configuration command. The BackboneFast feature is supported only when the switch is running PVST.

The switch tries to determine if it has an alternate path to the root switch. If the inferior BPDU arrives on a blocked port, the root port and other blocked ports on the switch become alternate paths to the root switch. (Self-looped ports are not considered alternate paths to the root switch.) If the inferior BPDU arrives on the root port, all blocked ports become alternate paths to the root switch. If the inferior BPDU arrives on the root port and there are no blocked ports, the switch assumes that it has lost connectivity to the root switch, causes the maximum aging time on the root port to expire, and becomes the root switch according to normal spanning-tree rules.

If the switch has alternate paths to the root switch, it uses these alternate paths to send a root link query (RLQ) request. The switch sends the RLQ request on all alternate paths to the root switch and waits for an RLQ reply from other switches in the network. If the switch determines

that it still has an alternate path to the root, it expires the maximum aging time on the port that received the inferior BPDU. If all the alternate paths to the root switch indicate that the switch has lost connectivity to the root switch, the switch expires the maximum aging time on the port that received the RLQ reply. If one or more alternate paths can still connect to the root switch, the switch makes all ports on which it received an inferior BPDU its designated ports and moves them from the blocking state (if they were in the blocking state), through the listening and learning states, and into the forwarding state.

You can enable BackboneFast to detect indirect link failures and to start the spanning-tree reconfiguration sooner.

Note If you use BackboneFast, you must enable it on all switches in the network. BackboneFast is not supported on Token Ring VLANs. This feature is supported for use with third-party switches.

The BackboneFast feature is supported only when the switch is running PVST.

Use the steps outlined in thefollowing table to enable BackboneFast:

Table 5-37: BackboneFast

Command	Purpose
spanning-tree backbonefast	Enable BackboneFast.

To disable the BackboneFast feature, use the **no spanning-tree backbonefast** global configuration command.

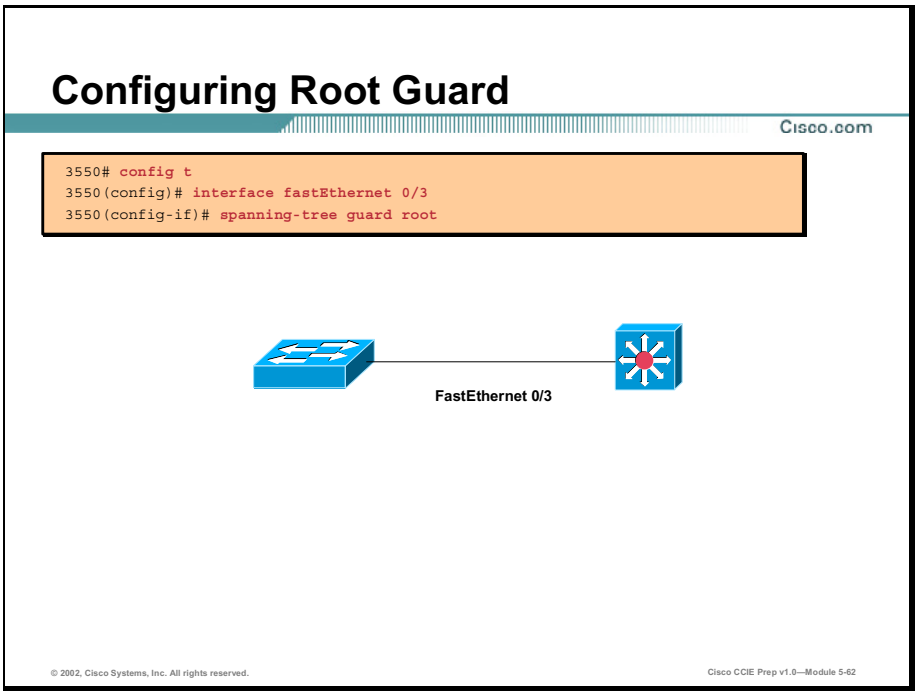

When a change in the spanning-tree topology occurs, a new root bridge is sometimes selected. If you let spanning-tree defaults dictate the election of the root bridge, you may end up with a non-preferred switch, such as an access layer switch, performing the root bridge function. You can avoid this situation by configuring root guard the switches in your network. If spanning-tree calculations cause an interface in the customer network to be selected as the root port, root guard then places the interface in the root-inconsistent (blocked) state to prevent the non-preferred switch from becoming the root switch or being in the path to the root.

Do not enable the root guard on interfaces to be used by the UplinkFast feature. With UplinkFast, the backup interfaces (in the blocked state) replace the root port in the case of a failure. However, if root guard is also enabled, all the backup interfaces used by the UplinkFast feature are placed in the root-inconsistent state (blocked) and are prevented from reaching the forwarding state.

Note	You cannot enable both root guard and loop guard at the same time.

Use the steps outlined in the table to enable root guard on an interface:

Table 5-38: Root Guard

Command	Purpose
interface interface-id	Enter interface configuration mode, and specify an interface to configure.
spanning-tree guard root	Enable root guard on the interface. By default, root guard is disabled on all interfaces.

To disable root guard, use the **no spanning-tree guard** interface configuration command.

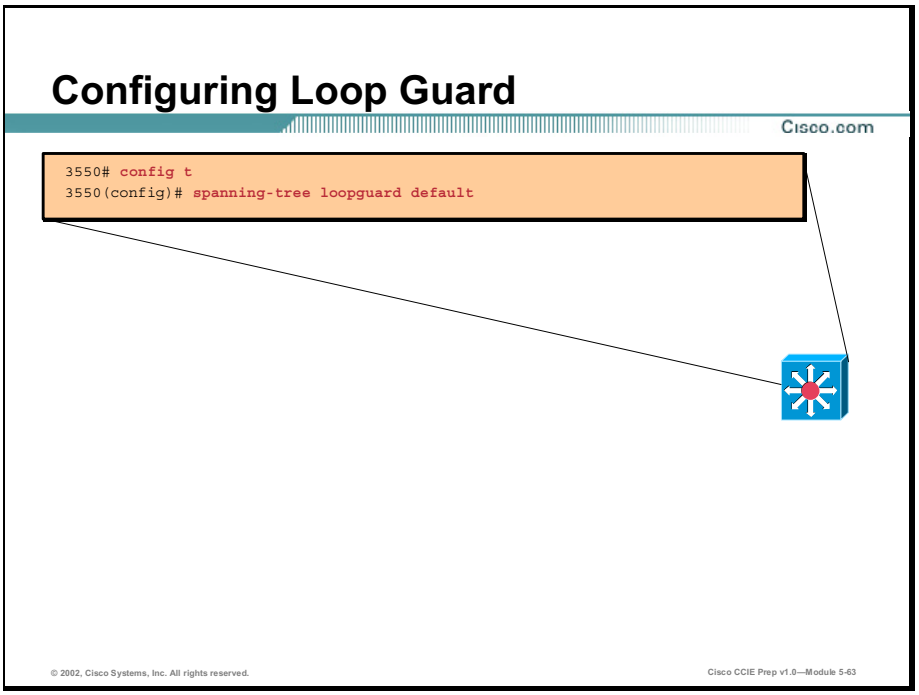

You can use loop guard to prevent alternate or root ports from becoming designated ports because of a failure that leads to a unidirectional link. This feature is most effective when it is configured on the entire switched network.

You can use loop guard to prevent alternate or root ports from becoming designated ports because of a failure that leads to a unidirectional link. This feature is most effective when it is configured on the entire switched network. Loop guard operates only on ports that are considered point-to-point by the spanning tree.

Note	You cannot enable both loop guard and root guard at the same time.

Use the steps outlined in the following table to enable loop guard:

Table 5-39: Root Guard

Command	Purpose
show spanning-tree active or **show spanning-tree mst**	Determine which ports are alternate or root ports.
spanning-tree loopguard default	Enable loop guard. By default, loop guard is disabled.

To globally disable loop guard, use the **no spanning-tree loopguard default** global configuration command. You can override the setting of the **no spanning-tree loopguard default** global configuration command by using the **spanning-tree guard loop** interface configuration command.

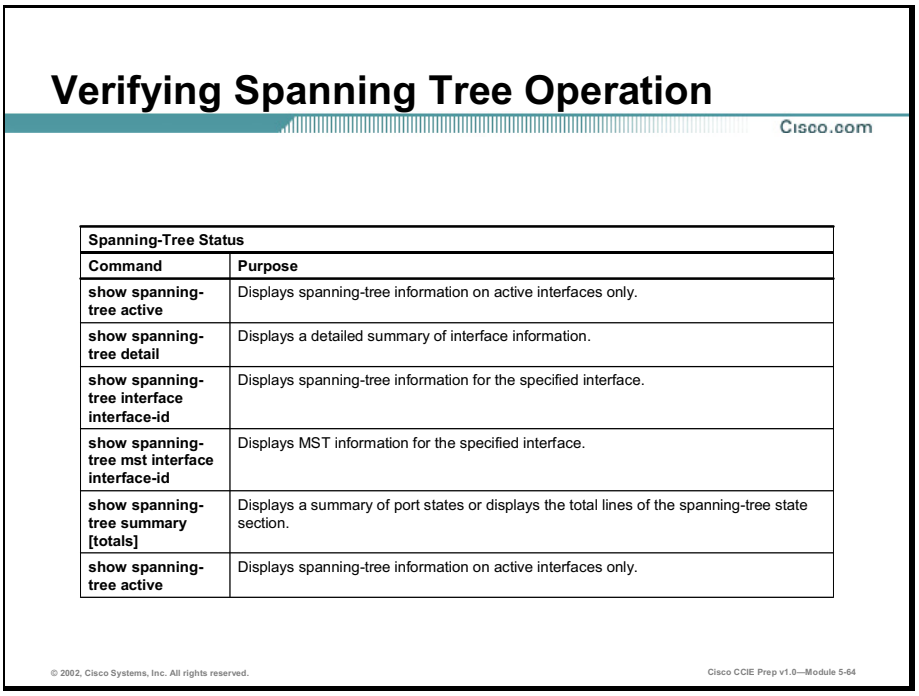

To display the spanning-tree status, use one or more of the following commands:

Table 5-40: Spanning-Tree Status

Command	Purpose
show spanning-tree active	Displays spanning-tree information on active interfaces only.
show spanning-tree detail	Displays a detailed summary of interface information.
show spanning-tree interface interface-id	Displays spanning-tree information for the specified interface.
show spanning-tree mst interface interface-id	Displays MST information for the specified interface.
show spanning-tree summary [totals]	Displays a summary of port states or displays the total lines of the spanning-tree state section.
show spanning-tree active	Displays spanning-tree information on active interfaces only.

For information about other keywords for the **show spanning-tree** privileged EXEC command, refer to the command reference for this release.

Monitoring and Analyzing Traffic

This section examines the use of SPAN and RSPAN to monitor and analyze traffic as it passes through the Catalyst 3550.

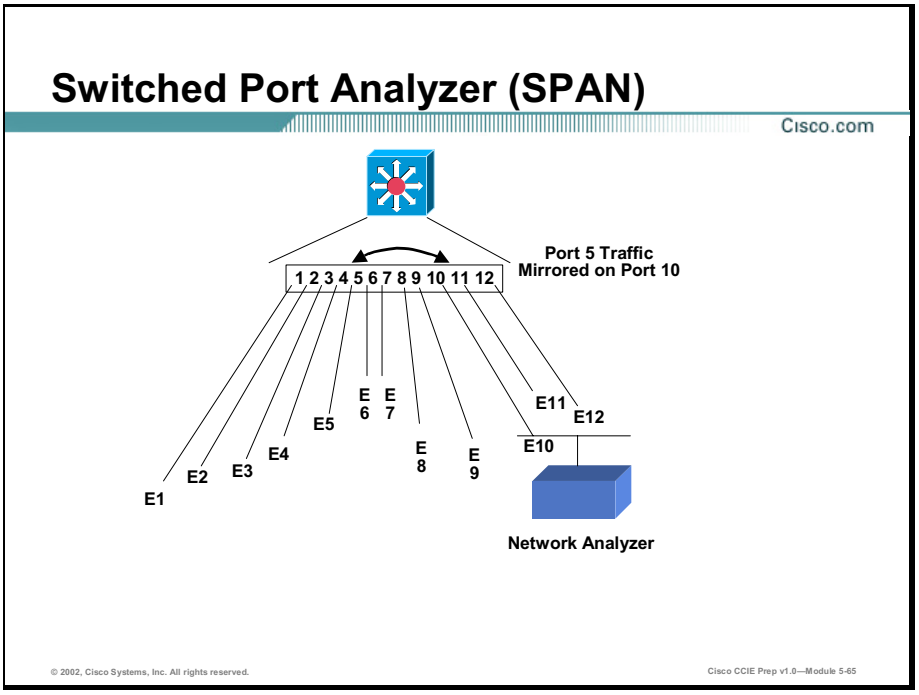

You can analyze network traffic passing through ports or VLANs by using SPAN to send a copy of the traffic to another port on the switch that has been connected to a SwitchProbe device or other Remote Monitoring (RMON) probe. SPAN mirrors received or sent (or both) traffic on a source port and received traffic on one or more source ports or source VLANs, to a destination port for analysis.

For example, in the figure above, all traffic on port 5 (the source port) is mirrored to port 10 (the destination port). A network analyzer on port 10 receives all network traffic from port 5 without being physically attached to port 5.

Only traffic that enters or leaves source ports or traffic that enters source VLANs can be monitored by using SPAN; traffic that gets routed to ingress source ports or source VLANs cannot be monitored. For example, if incoming traffic is being monitored, traffic that gets routed from another VLAN to the source VLAN is not monitored; however, traffic that is received on the source VLAN and routed to another VLAN is monitored.

Creating a SPAN Session and Specifying Ports to Monitor

Beginning in privileged EXEC mode, follow these steps to create a SPAN session and specify the source (monitored) and destination (monitoring) ports:

Table 5-41: SPAN Session

Command	Purpose
no monitor session {*session_number* | **all | local | remote**}	Clear any existing SPAN configuration for the session. For *session_number*, specify 1 or 2. Specify **all** to remove all SPAN sessions, **local** to remove all local sessions, or **remote** to remove all remote SPAN sessions.
monitor session *session_number* **source interface** *interface-id* [, | -] [both | rx | tx]	Specify the SPAN session and the source port (monitored port). For *session_number*, specify 1 or 2. For *interface-id*, specify the source port to monitor. Valid interfaces include physical interfaces and port-channel logical interfaces (**port-channel** *port-channel-number*). (Optional) **[, | -]** Specify a series or range of interfaces. Enter a space after the comma; enter a space before and after the hyphen. (Optional) Specify the direction of traffic to monitor. If you do not specify a traffic direction, the source interface sends both sent and received traffic. Only received (rx) traffic can be monitored on additional source ports. ■ **both**—Monitor both received and sent traffic. ■ **rx**—Monitor received traffic. ■ **tx**—Monitor sent traffic.
monitor session *session_number* **destination interface** *interface-id* [**encapsulation** {**dot1q | isl**}]	Specify the SPAN session and the destination port (monitoring port). For *session_number*, specify 1 or 2. For *interface-id*, specify the destination port. Valid interfaces include physical interfaces. (Optional) Specify the encapsulation header for outgoing packets. If not specified, packets are sent in native form. **isl**—Use ISL encapsulation. **dot1q**—Use 802.1Q encapsulation.

Removing Ports from a SPAN Session

Beginning in privileged EXEC mode, follow these steps to remove a port as a SPAN source for a session:

Table 5-42: SPAN Source

Command	Purpose
no monitor session *session_number* **source interface** *interface-id* **[,** | **-]** [**both** | **rx** | **tx**]	Specify the characteristics of the source port (monitored port) and SPAN session to remove. For *session*, specify 1 or 2. For *interface-id*, specify the source port to no longer monitor. Valid interfaces include physical interfaces and port-channel logical interfaces (**port-channel** *port-channel-number*). (Optional) Use **[, | -]** to specify a series or range of interfaces if they were configured. This option is valid when monitoring only received traffic. Enter a space after the comma; enter a space before and after the hyphen. (Optional) Specify the direction of traffic (**both, rx**, or **tx**) to no longer monitor. If you do not specify a traffic direction, both transmit and receive are disabled.

To remove a source or destination port from the SPAN session, use the **no monitor session** *session_number* **source interface** *interface-id* global configuration command or the **no monitor session** *session_number* **destination interface** *interface-id* global configuration

command. To change the encapsulation type back to the default (native), use the **monitor session** *session_number* **destination interface** *interface-id* without the **encapsulation** keyword.

Specifying VLANs to Monitor

VLAN monitoring is similar to port monitoring. Beginning in privileged EXEC mode, follow these steps to specify VLANs to monitor:

Table 5-43: VLANs to Monitor

Command	Purpose
no monitor session {*session_number* | **all** | **local** | **remote**}	Clear any existing SPAN configuration for the session. For *session_number*, specify 1 or 2. Specify **all** to remove all SPAN sessions, **local** to remove all local sessions, or **remote** to remove all remote SPAN sessions.
monitor session ***session_number*** **source vlan** ***vlan-id*** [, | -] **rx**	Specify the SPAN session and the source VLANs (monitored VLANs). You can monitor only received (**rx**) traffic on VLANs. For *session_number*, specify 1 or 2. For *vlan-id*, the range is 1 to 4094; do not enter leading zeros. (Optional) Use a comma (**,**) to specify a series of VLANs, or use a hyphen (**-**) to specify a range of VLANs. Enter a space after the comma; enter a space before and after the hyphen.
monitor session ***session_number*** **destination interface** ***interface-id*** **[encapsulation** {**dot1q** | **isl**}]	Specify the SPAN session and the destination port (monitoring port). For *session_number*, specify 1 or 2. For *interface-id*, specify the destination port. Valid interfaces include physical interfaces. (Optional) Specify the encapsulation header for outgoing packets. If not specified, packets are sent in native form. ■ **isl**—Use ISL encapsulation. ■ **dot1q**—Use 802.1Q encapsulation.

To remove one or more source VLANs or destination ports from the SPAN session, use the **no monitor session** *session_number* **source vlan** *vlan-id* **rx** global configuration command or the **no monitor session** *session_number* **destination interface** *interface-id* global configuration command.

Specifying VLANs to Filter

Beginning in privileged EXEC mode, follow these steps to limit SPAN source traffic to specific VLANs:

Table 5-44: Limit SPAN Source Traffic

Command	Purpose
no monitor session {*session_number* | **all** | **local** | **remote**}	Clear any existing SPAN configuration for the session. For *session_number*, specify 1 or 2. Specify **all** to remove all SPAN sessions, **local** to remove all local sessions, or **remote** to remove all remote SPAN sessions.
monitor session session_number **source interface** interface-id **rx**	Specify the characteristics of the source port (monitored port) and SPAN session. For *session_number*, specify 1 or 2. For *interface-id*, specify the source port to monitor. The interface specified must already be configured as a trunk port.
monitor session session_number **filter vlan** vlan-id *[, | -]*	Limit the SPAN source traffic to specific VLANs. For *session_number*, specify 1 or 2. For *vlan-id*, the range is 1 to 4094; do not enter leading zeros. (Optional) Use a comma (**,**) to specify a series of VLANs or use a hyphen (**-**) to specify a range of VLANs. Enter a space after the comma; enter a space before and after the hyphen.
monitor session session_number **destination interface** *interface-id*	Specify the characteristics of the destination port (monitoring port) and SPAN session. For *session_number*, specify 1 or 2. For *interface-id*, specify the destination port. Valid interfaces include physical interfaces.

To monitor all VLANs on the trunk port, use the **no monitor session** *session_number* **filter** global configuration command.

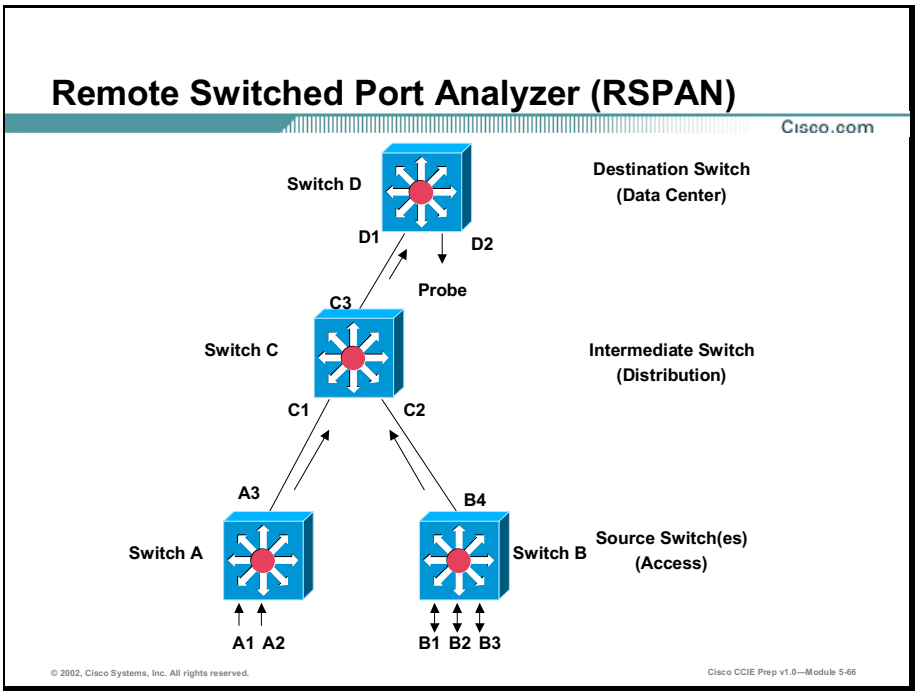

RSPAN extends SPAN by enabling remote monitoring of multiple switches across your network. The traffic for each RSPAN session is carried over a user-specified RSPAN VLAN that is dedicated for that RSPAN session in all participating switches. The SPAN traffic from the sources is copied onto the RSPAN VLAN through a reflector port and then forwarded over trunk ports that are carrying the RSPAN VLAN to any RSPAN destination sessions monitoring the RSPAN VLAN, as shown in the figure above.

First create an RSPAN VLAN that *does not* exist for the RSPAN session in any of the switches that will participate in RSPAN. With VTP enabled in the network, you can create the RSPAN VLAN in one switch, and VTP propagates it to the other switches in the VTP domain for VLAN-IDs that are lower than 1005.

Use VTP pruning to get efficient flow of RSPAN traffic, or manually delete the RSPAN VLAN from all trunks that do not need to carry the RSPAN traffic.

After creating the RSPAN VLAN, begin in privileged EXEC mode, and follow these steps to start an RSPAN source session and to specify the source (monitored) ports and the destination RSPAN VLAN.

Table 5-45: RSPAN

Command	Purpose
no monitor session {*session_number* | **all** | **local** | **remote**}	Clear any existing RSPAN configuration for the session. For *session_number*, specify 1 or 2. Specify **all** to remove all RSPAN sessions, **local** to remove all local sessions, or **remote** to remove all remote SPAN sessions.
monitor session *session_number* **source interface** *interface-id* [, | -] [**both** | **rx** | **tx**]	Specify the RSPAN session and the source port (monitored port). For *session_number*, specify 1 or 2. For *interface-id*, specify the source port to monitor. Valid interfaces include physical interfaces and port-channel logical interfaces (**port-channel** *port-channel-number*). (Optional) [**, | -**] Specify a series or range of interfaces. Enter a space after the comma; enter a space before and after the hyphen. (Optional) Specify the direction of traffic to monitor. If you do not specify a traffic direction, the source interface sends both sent and received traffic. Only received (**rx**) traffic can be monitored on additional source ports. ■ **both**—Monitor both received and sent traffic. ■ **rx**—Monitor received traffic. ■ **tx**—Monitor sent traffic.
monitor session *session_number* **destination remote vlan** *vlan-id* **reflector-port** *interface*	Specify the RSPAN session, the destination remote VLAN, and the reflector port. For *session_number*, enter 1 or 2. For *vlan-id*, specify the RSPAN VLAN to carry the monitored traffic to the destination port. For *interface*, specify the interface that will flood the RSPAN traffic onto the RSPAN VLAN.

Creating an RSPAN Destination Session

Beginning in privileged EXEC mode, follow these steps to create an RSPAN destination session and to specify the source RSPAN VLAN and the destination port:

Table 5-46: RSPAN VLAN

Command	Purpose
monitor session *session_number* **source remote vlan** *vlan-id*	Specify the RSPAN session and the source RSPAN VLAN. For *session_number*, specify 1 or 2. For *vlan-id*, specify the source RSPAN VLAN to monitor.
monitor session *session_number* **destination interface** *interface-id* [**encapsulation** {**dot1q** | **isl**}]	Specify the RSPAN session and the destination interface. For *session_number*, specify 1 or 2. For *interface-id*, specify the destination interface. (Optional) Specify the encapsulation header for outgoing packets. If not specified, packets are sent in native form. ■ **isl**—Use ISL encapsulation. ■ **dot1q**—Use 802.1Q encapsulation.

Removing Ports from an RSPAN Session

Beginning in privileged EXEC mode, follow these steps to remove a port as an RSPAN source for a session:

Table 5-47: RSPAN

Command	Purpose
no monitor session *session_number* **source interface** *interface-id* **[, | -] [both | rx | tx]**	Specify the characteristics of the RSPAN source port (monitored port) to remove. For *session_number*, specify 1 or 2. For *interface-id*, specify the source port to no longer monitor. Valid interfaces include physical interfaces and port-channel logical interfaces (**port-channel** *port-channel-number*). (Optional) Use **[, | -]** to specify a series or range of interfaces if they were configured. Enter a space after the comma; enter a space before and after the hyphen. (Optional) Specify the direction of traffic (**both, rx**, or **tx**) to no longer monitor. If you do not specify a traffic direction, both transmit and receive are disabled.

Specifying VLANs to Monitor

VLAN monitoring is similar to port monitoring. Beginning in privileged EXEC mode, follow these steps to specify VLANs to monitor:

Table 5-47: VLANs to Monitor

Command	Purpose
no monitor session {*session_number* | **all** | **local** | **remote**}	Clear any existing SPAN configuration for the session. For *session_number*, specify 1 or 2. Specify **all** to remove all SPAN sessions, **local** to remove all local sessions, or **remote** to remove all remote SPAN sessions.
monitor session *session_number* **source vlan** *vlan-id* **[, | -] rx**	Specify the RSPAN session and the source VLANs (monitored VLANs). You can monitor only received (**rx**) traffic on VLANs. For *session_number*, specify 1 or 2. For *vlan-id*, the range is 1 to 4094; do not enter leading zeros. (Optional) Use a comma (**,**) to specify a series of VLANs, or use a hyphen (**-**) to specify a range of VLANs. Enter a space after the comma; enter a space before and after the hyphen.
monitor session *session_number* **destination remote vlan** *vlan-id* **reflector port** *interface*	Specify the RSPAN session, the destination remote VLAN, and the reflector port. For *session_number*, enter 1 or 2. For *vlan-id*, specify the RSPAN VLAN to carry the monitored traffic to the destination port. For *interface*, specify the interface that will flood the RSPAN traffic to the RSPAN VLAN.

To remove one or more source VLANs from the RSPAN session, use the **no monitor session** *session_number* **source vlan** *vlan-id* **rx** global configuration command.

Specifying VLANs to Filter

Beginning in privileged EXEC mode, follow these steps to limit RSPAN source traffic to specific VLANs:

Table 5-48: VLANs to Filter

Command	Purpose
no monitor session {*session_number* | **all** | **local** | **remote**}	Clear any existing SPAN configuration for the session. For *session_number*, specify 1 or 2. Specify **all** to remove all SPAN sessions, **local** to remove all local sessions, or **remote** to remove all remote SPAN sessions.
monitor session *session_number* **source interface** *interface-id* **rx**	Specify the characteristics of the source port (monitored port) and RSPAN session. For *session_number*, specify 1 or 2. For *interface-id*, specify the source port to monitor. The interface specified must already be configured as a trunk port.
monitor session *session_number* **filter vlan** *vlan-id* [, | -]	Limit the RSPAN source traffic to specific VLANs. For *session_number*, specify 1 or 2. For *vlan-id*, the range is 1 to 4094; do not enter leading zeros. (Optional) Use a comma (**,**) to specify a series of VLANs or use a hyphen (**-**) to specify a range of VLANs. Enter a space after the comma; enter a space before and after the hyphen.
monitor session *session_number* **destination remote vlan** *vlan-id* **reflector port** *interface*	Specify the RSPAN session, the destination remote VLAN, and the reflector port. For *session_number*, enter 1 or 2. For *vlan-id*, specify the RSPAN VLAN to carry the monitored traffic to the destination port. For *interface*, specify the interface that will flood the RSPAN traffic to the RSPAN VLAN.

Verifying SPAN and RSPAN

Cisco.com

```
Switch# show monitor session 1
Session 1
---------
Type: Remote Source Session
Source Ports:
    RX Only:                Fa0/3
    TX Only:       None
    Both:          None
Source VLANs:
    RX Only:       None
    TX Only:       None
    Both:          None
Source RSPAN VLAN: None
Destination Ports: None
          Encapsulation: Native
Reflector Port:      Fa0/4
Filter VLANs:      None
Dest RSPAN VLAN:    901
```

To display the status of the current SPAN or RSPAN configuration, use the **show monitor** privileged EXEC command.

Fallback Bridging

This section describes how to configure fallback bridging (VLAN bridging) on your switch. With fallback bridging, you can forward non-IP packets that the switch does not route between VLAN bridge domains and routed ports. To use this feature, you must have the enhanced multilayer software (EMI) image installed on your switch.

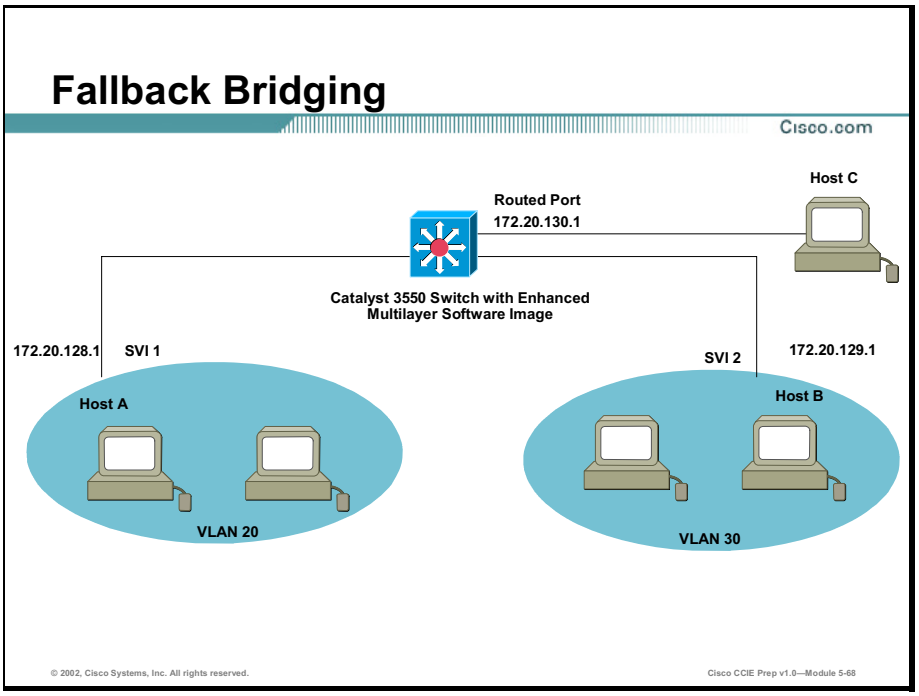

With fallback bridging, the switch bridges together two or more VLANs or routed ports, essentially connecting multiple VLANs within one bridge domain. Fallback bridging forwards traffic that the switch does not route and forwards traffic belonging to a nonroutable protocol such as DECnet.

Fallback bridging does not allow the spanning trees from the VLANs being bridged to collapse; each VLAN has its own spanning-tree instance and a separate spanning tree, called the VLAN-bridge spanning tree, which runs on top of the bridge group to prevent loops.

A VLAN bridge domain is represented with switch virtual interface (SVI). A set of SVIs and routed ports (which do not have any VLANs associated with them) can be configured (grouped together) to form a bridge group. Recall that an SVI represents a VLAN of switch ports as one interface to the routing or bridging function in the system. You associate only one SVI with a VLAN, and you configure an SVI for a VLAN only when you want to route between VLANs, to fallback-bridge nonroutable protocols between VLANs, or to provide IP host connectivity to the switch. A routed port is a physical port that acts like a port on a router, but it is not connected to a router. A routed port is not associated with a particular VLAN, does not support VLAN subinterfaces, but behaves like a normal routed interface.

A bridge group is an internal organization of network interfaces on a switch. Bridge groups cannot be used to identify traffic switched within the bridge group outside the switch on which they are defined. Bridge groups on the same switch function as distinct bridges; that is, bridged

traffic and bridge protocol data units (BPDUs) are not exchanged between different bridge groups on a switch. An interface can be a member of only one bridge group. Use a bridge group for each separately bridged (topologically distinct) network connected to the switch.

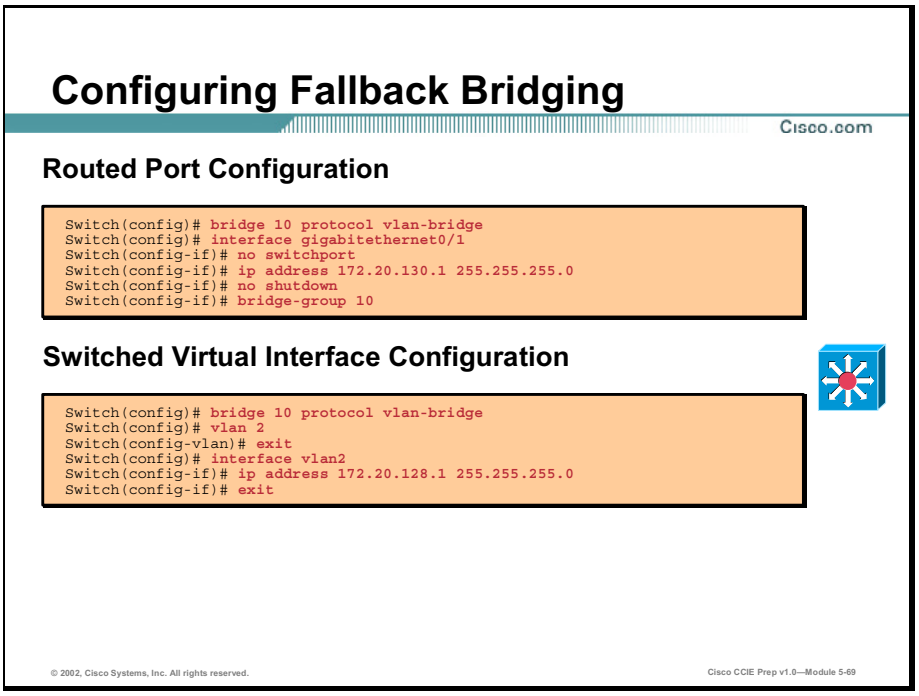

To configure fallback bridging for a set of SVIs or routed ports, these interfaces must be assigned to bridge groups. All interfaces in the same group belong to the same bridge domain. Each SVI or routed port can be assigned to only one bridge group. A maximum of 31 bridge groups can be configured on the switch.

Note	The protected port feature is not compatible with fallback bridging. When fallback bridging is enabled, it is possible for packets to be forwarded from one protected port on a switch to another protected port on the same switch if the ports are in different VLANs.

Use the steps outlined in the following table to create a bridge group and assign an interface to it:

Table 5-49: Bridge Group

Command	Purpose
bridge *bridge-group* **protocol vlan-bridge**	Assign a bridge group number, and specify the VLAN-bridge spanning-tree protocol to run in the bridge group. The **ibm** and **dec** keywords are not supported. For *bridge-group*, specify the bridge group number. The range is 1 to 255. You can create up to 31 bridge groups. Frames are bridged only among interfaces in the same group..
interface *interface-id*	Enter interface configuration mode, and specify the interface on which you want to assign the bridge group. The specified interface must be one of these: ■ A routed port: a physical port that you have configured as a Layer 3 port by entering the **no switchport** interface configuration command. ■ An SVI: a VLAN interface that you created by using the **interface vlan** *vlan-id* global configuration command. These ports must have IP addresses assigned to them.
bridge-group *bridge-group*	Assign the interface to the bridge group created in Step 2. By default, the interface is not assigned to any bridge group. An interface can be assigned to only one bridge group.

To remove a bridge group, use the **no bridge** *bridge-group* global configuration command. To remove an interface from a bridge group and to remove the bridge group, use the **no bridge-group** *bridge-group* interface configuration command.

Summary

This section summarizes the key points discussed in this lesson.

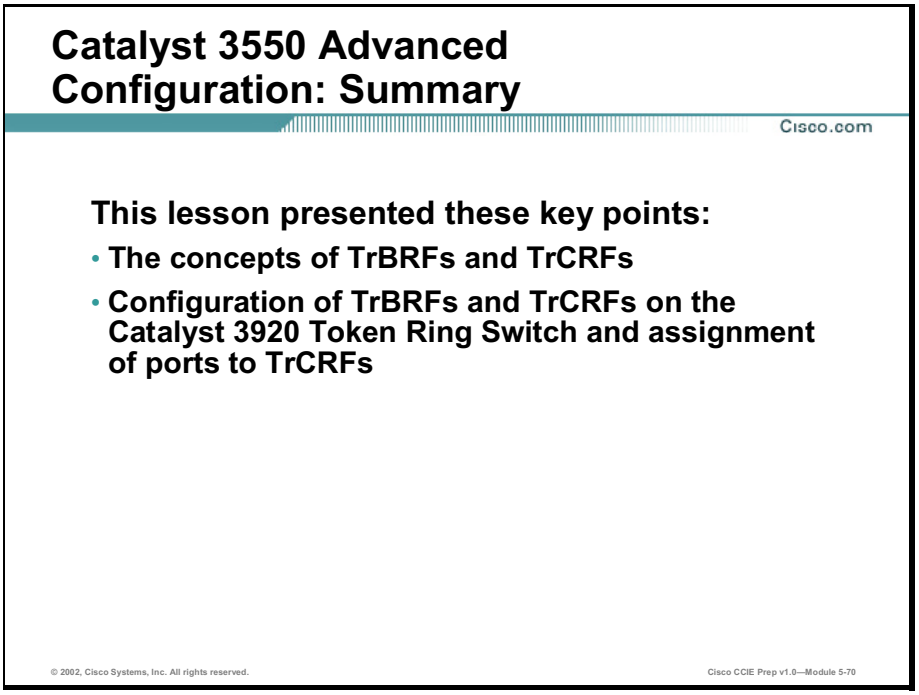

Next Steps

After completing this lesson, go to:

- Distance-Vector Routing Protocols

References

For additional information, refer to these resources:

- CCIE Professional Development : Cisco Lan Switching by Kennedy Clark

Lesson Assessment (Quiz)

Q1) Which of the following features shut down a PortFast enabled port when a BPDU is received on that port?

A) RootGuard

B) BPDUGuard

C) LoopGuard

D) 802.1X Guard

E) PAgPGuard

Q2) _____ extends SPAN by enabling remote monitoring of multiple switches across your network.

A) SwitchProbe

B) RMON

C) RSPAN

D) Extended SPAN

With _____ you can forward non-IP packets that the switch does not route between VLAN bridge domains and routed ports.

E) 802.1Q Tunneling

F) Layer 2 Protocol Tunneling

G) InterVLAN routing

H) Fallback Bridging

5

Distance-Vector Routing Protocols

Overview

This module briefly examines routing protocols in general, followed by a review of the major distance-vector routing protocols, Routing Information Protocol (RIP), RIPv2, and Enhanced Interior Gateway Routing Protocol (EIGRP).

Upon completing this module, you will be able to:

- Explain the various fields used in the routing table
- List the major differences between Link-State routing protocols and distance-vector routing protocols
- Perform advanced configurations of RIP, RIPv2, and EIGRP

Outline

The module contains these lessons:

- Routing Information Protocol (RIP)
- Enhanced Interior Gateway Routing Protocol (EIGRP)

Routing Information Protocol (RIP)

Overview

This lesson will cover the basic and advanced configuration of Routing Information Protocol (RIP). This lesson also covers the operation and tuning of RIPv1 and RIPv2.

Importance

RIP performs its job well, and is a very popular routing protocol. RIP has evolved over the years from a classful routing protocol, RIP version 1 (RIPv1), to a classless routing protocol, RIP version 2 (RIPv2). RIP is one of the routing protocols tested on the Cisco Certified Internetworking Expert (CCIE) Lab.

Objectives

Upon completing this lesson, you will be able to:

- Describe how RIP operates as a routing protocol
- Perform advanced configurations of RIPv1 and RIPv2

Learner Skills and Knowledge

To fully benefit from this lesson, you must have these prerequisite skills and knowledge:

- Passed the CCIE written qualification exam
- Solid understanding of Internet Protocol (IP) addressing and Cisco router fundamentals

Outline

This lesson includes these sections:

- Overview
- RIP
- RIP Version 2 (RIPv2)
- Optional RIP Configuration Tasks
- Trouble Shooting
- Summary
- Lesson Assessment (Quiz)

RIP

This section provides an overview of RIP and describes how to configure it on a Cisco router.

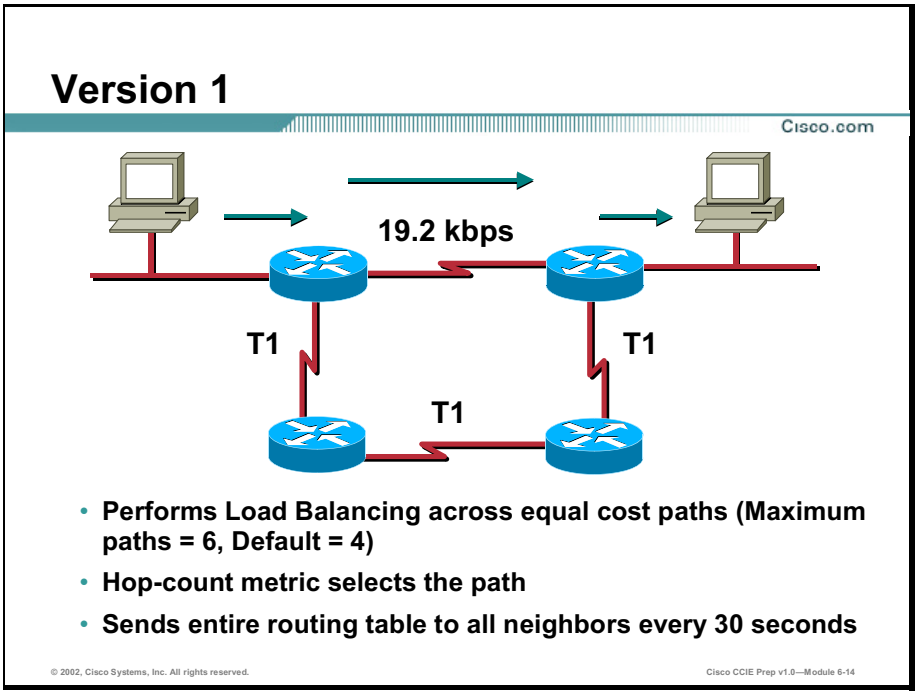

RIP version 1 (RIPv1) is described in RFC 1058, and an enhanced version, RIP version 2 (RIPv2), a classless routing protocol, is defined in Requests for Comment (RFC) 1721, 1722, and 1723.

Key characteristics of RIP include the following:

- It is a distance-vector routing protocol that operates on User Datagram Protocol (UDP) port 520.
- Hop count is used as the metric for path selection.
- It has an administrative distance of 120.
- The maximum allowable hop count is 15.
- Routing updates are broadcast every 30 seconds by default.
- RIPv1 is a classful routing protocol. Classful routing protocols do not send the subnet mask along with the routing update.
- RIPv2 is a classless routing protocol. Classless routing protocols do send the subnet mask along with the routing update.

- RIP is capable of load balancing over 6 equal cost paths (four is the default).
- Defining the maximum number of parallel paths allowed in a routing table enables RIP load balancing. With RIP, the paths must be equal-cost paths. If the maximum number of paths command is set to 1 (one), load balancing is disabled.
- Since RIP is a distance-vector routing protocol, it is vulnerable to split horizon issues, especially in Non-Broadcast Multi-Access (NBMA) networks, such as Frame Relay.

RIP Version 2 (RIPv2)

RIPv2 is not a new routing protocol, but an extension of RIPv1 provided by RFC 1721, 1722, and 1723.

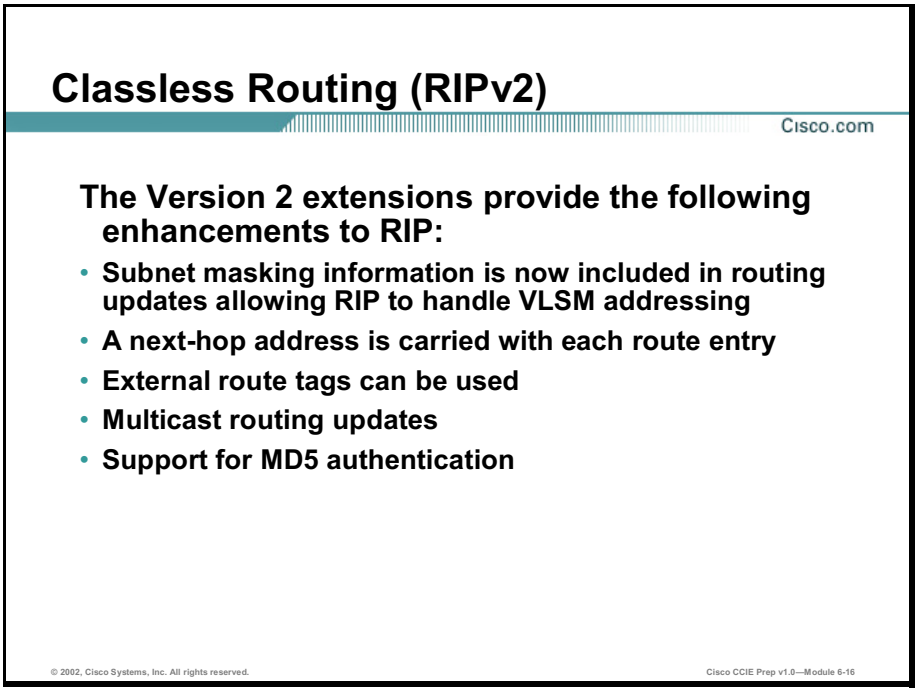

The most significant of all the enhancements is the support for Variable Length Subnet Mask (VLSM), making RIPv2 a classless routing protocol.

Most of RIPv2's operational procedures and timers are identical to RIPv1. However, RIPv2 uses the multicast address of 224.0.0.9 to send updates versus the general all hosts broadcast used by RIPv1.

RIPv2 is fully backward compatible with RIPv1. This is accomplished by means of a compatibility switch and a receive control switch, as defined in RFC 1723. Essentially, these switches allow you to control what type of RIP updates the router sends and receives. The router can be configured to receive only Version 1 updates, only Version 2 updates, both, or none. The router can send only Version 1 updates, send Version 2 updates as a broadcast message, send Version 2 updates as a multicast, or send no updates at all. The switches can be manually set with the following interface command:

```
Router(config-if)# ip rip [send | receive] version [1 | 2 | 1 2]
```

To enable RIPv2 globally use the **version 2** command, from the router configuration mode:

A) Router(config-router)# **version 2**

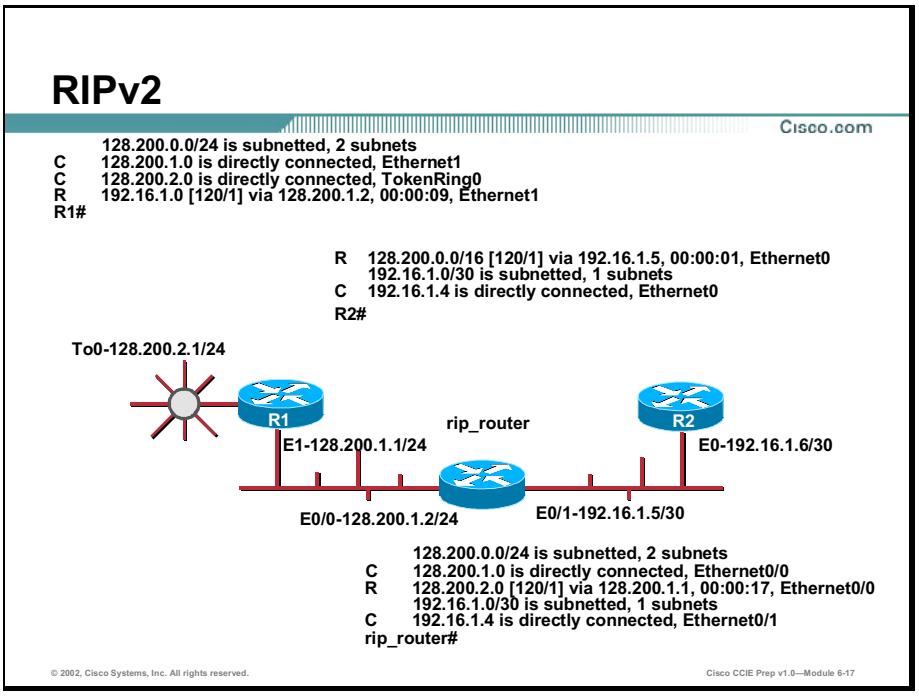

In this example, the Token Ring network on R1 is configured to send and receive both RIPv1 and RIPv2 updates. The Ethernet segment off of R1, however, will send and receive only Version 2 updates. The rip_router and R2 are configured to send and receive only RIPv2 updates.

RIPv2 Configuration

```
hostname R1
!
interface Ethernet1
 ip address 128.200.1.1 255.255.255.0
 ip rip send version 2
 ip rip receive version 2
 media-type 10BaseT
!
interface TokenRing0
 ip address 128.200.2.1 255.255.255.0
 ip rip send version 1 2
 ip rip receive version 1 2
 ring-speed 16
!
router rip
 version 2
 network 128.200.0.0
 no auto-summary
```



```
hostname rip_router
!
interface Ethernet0/0
 ip address 128.200.1.2 255.255.255.0
 ip rip send version 2
 ip rip receive version 2
!
interface Ethernet0/1
 ip address 192.16.1.5 255.255.255.252
 ip rip send version 2
 ip rip receive version 2
!
router rip
 version 2
 network 128.200.0.0
 network 192.16.1.0
 no auto-summary


hostname R2
!
interface Ethernet0
 ip address 192.16.1.6 255.255.255.252
 ip rip send version 2
 ip rip receive version 2
!
router rip
 version 2
 network 192.16.1.0
 no auto-summary
```

B)

Optional RIP Configuration Tasks

This section covers the commands to perform optional configurations tasks in RIP, such as modifying RIP timers, setting the maximum number of paths to load balance across, and controlling RIP update traffic.

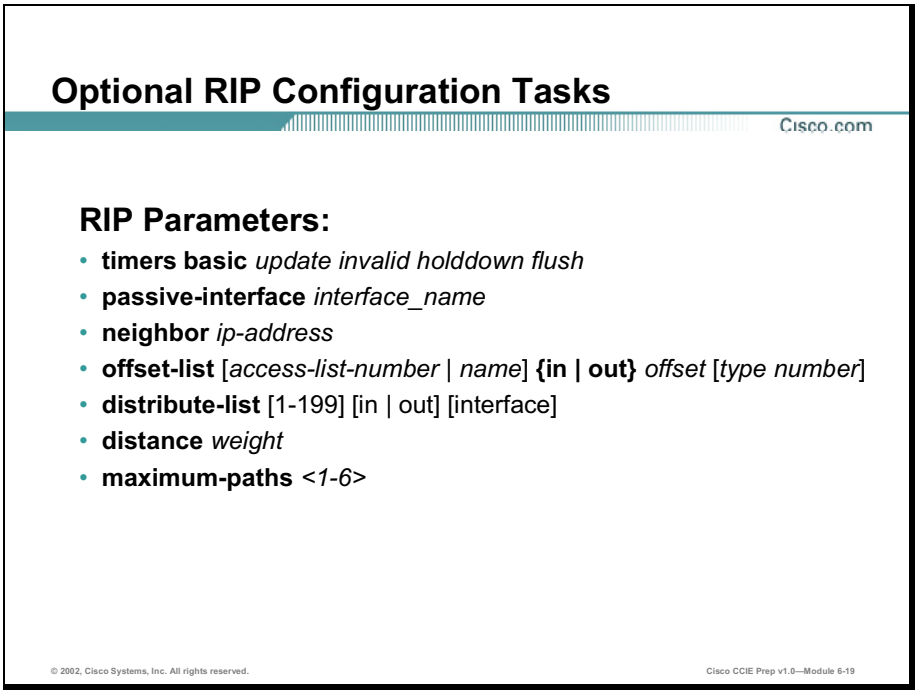

The following is a list of some of the common adjustable parameters within RIP.

- **timers basic** *update invalid holddown flush* This allows the user to set the update, invalid, holddown, and flush timers for RIP.
- **passive-interface** *interface_name* This command prevents the sending of routing updates on an interface; however, the router still listens to updates received from that interface.
- **neighbor** *ip-address* This command defines a RIP neighbor to exchange unicast updates with and should be used in conjunction with the **passive-interface** command.
- **offset-list** [*access-list-number* | *name*] {**in** | **out**} *offset* [*type number*] Use this command to increase the value of the routing metrics. Default values for the *access-list-number* argument are 0-99. The metric offset cannot exceed 16.
- **distribute-list** [**1-199**] [**in** | **out**] [**interface**] Use this command to call a standard or extended access list to filter inbound and/or outbound routing updates.
- **distance** *weight* {*ip-address* {*ip-address mask*}} [*ip standard list] [ip extended list*] Use this command to change the administrative distance of routes received from a neighbor. If

the IP address and **wildcard_mask** are omitted, all routes for that protocol will be set to the distance value.

- **maximum-paths** <*1-6*> Use this command to configure the maximum number of equal cost paths to load balance across. The default setting is 4. A setting of 1 disables load balancing.

Troubleshooting

This section discusses the major troubleshooting commands for RIP.

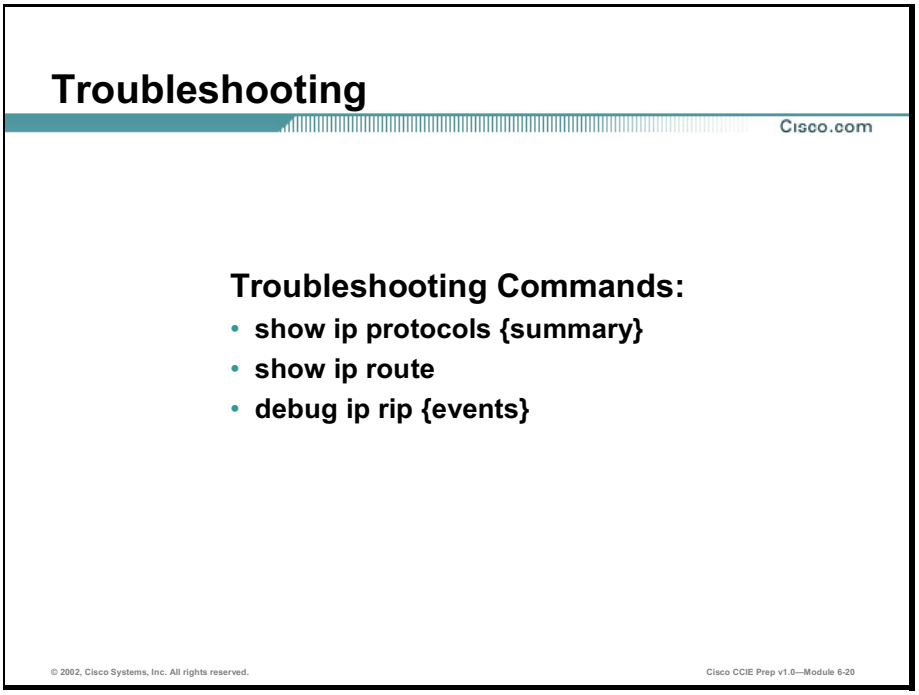

show ip protocols {summary} Command

This command displays all routing protocols, detailed timer and metric information, as well as routing update information.

```
R20# show ip protocols
Routing Protocol is "rip"  ←Routing Protocol Type
  Sending updates every 30 seconds, next due in 29 seconds
  Invalid after 180 seconds, hold down 180, flushed after 240    ←Timer
information
  Outgoing update filter list for all interfaces is        ←Distribute list (if
any)
  Incoming update filter list for all interfaces is
  Default redistribution metric is 2               ←Default metric
  Redistributing: rip, eigrp 2001               ←Redistribution is on
  Default version control: send version 1, receive any version
    Interface        Send  Recv   Key-chain
    Ethernet0/0      1     1 2    ←RIP Versions running
  Routing for Networks:        ←Networks participating in RIP
    128.200.0.0
  Passive Interface(s):
    Ethernet0/1            ←Network listening to RIP
  Routing Information Sources:
    Gateway         Distance      Last Update
```

```
    128.200.1.1          120      00:00:07 ←RIP Neighbors
  Distance: (default is 120)        ←Administrative Distance
```

show ip route Command

This command lists the router's current routing table, and the one on which it makes forwarding decisions. It is possible for a route to exist, or be known to the router, but only the routes with the shortest administrative distances are listed. The output from this command lists what routing protocol the route is from; in the case of the example, R for RIP. The numbers in the bracket behind the route is the administrative distance of the route followed by the hop count. The via field explains who the route is from, how long ago an update was received, and by what interface.

```
R20 show ip route
Gateway of last resort is not set

     128.200.0.0/16 is variably subnetted, 4 subnets, 2 masks
R       128.200.10.0/24 [120/1] via 128.200.1.1, 00:00:17, Ethernet0/0
C       128.200.1.0/24 is directly connected, Ethernet0/0
R       128.200.2.0/24 [120/1] via 128.200.1.1, 00:00:17, Ethernet0/0
C       128.200.3.16/29 is directly connected, Ethernet0/1
```

In this instance, the route 128.200.10.0/24 has a metric of 120, and is one hop away. The RIP neighbor providing information about the route is 128.200.1.1, and it sent the last update 17 seconds ago. This is also the next-hop for the targeted network in the routing table. R20 received it through its Ethernet 0/0 port. This is the next-hop interface to reach the destination network.

debug ip rip {events} Command

This command shows all the RIP activity occurring in the router and also displays exactly which interfaces are advertising and receiving routes. The RIP version of the update is also displayed, along with the metric of each route in the update.

```
R21# debug ip rip
1d02h: RIP: received v1 update from 128.200.10.2 on TokenRing1
1d02h:      128.200.10.0 in 1 hops
1d02h: RIP: sending v1 update to 255.255.255.255 via Ethernet1 (128.200.1.1)
1d02h:      subnet  128.200.10.0, metric 1
1d02h:      subnet  128.200.2.0, metric 1
1d02h: RIP: sending v1 update to 255.255.255.255 via TokenRing0 (128.200.2.1)
1d02h:      subnet  128.200.10.0, metric 1
1d02h:      subnet  128.200.1.0, metric 1
1d02h: RIP: sending v1 update to 128.200.10.2 via TokenRing1 (128.200.10.1)
1d02h:      subnet  128.200.10.0, metric 1
1d02h:      subnet  128.200.1.0, metric 1
1d02h:      subnet  128.200.2.0, metric 1
```

Summary

This section summarizes the key points discussed in this lesson.

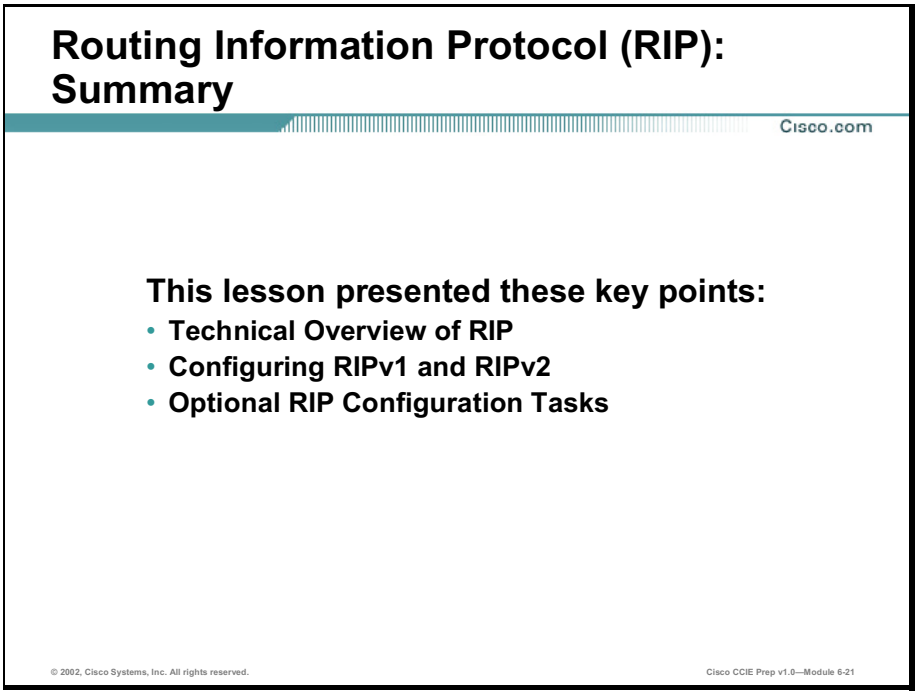

Next Steps

After completing this lesson, go to:

- Enhanced Interior Gateway Routing Protocol (EIGRP)

References

For additional information, refer to these resources:

- *Routing TCP/IP Volume I* by Jeff Doyle

Lesson Assessment (Quiz)

Q1) If the router receives a route update from a RIP neighbor and an internal BGP neighbor for the same route, which one is more believable?

Q2) If RIP has the passive interface command enabled for an interface, will RIP receive RIP routes on that interface? (Assume there is a downstream RIP device.)

A) Yes

B) No

Q3) What protocol and port number does RIPv2 use for communication with its RIP neighbors?

A) TCP 500

B) UDP 500

C) TCP 88

D) None of the above

Enhanced Interior Gateway Routing Protocol (EIGRP)

Overview

This lesson will examine the concepts and configuration of Cisco's Enhanced Interior Gateway Routing Protocol (EIGRP). This lesson also covers the operation and tuning of EIGRP, including performing manual route summarization to reduce the scope of EIGRP queries.

Importance

EIGRP is a Cisco proprietary Interior Gateway Protocol (IGP). Cisco Certified Internetworking Expert (CCIE) candidates should know how to configure EIGRP in both a Local Area Network (LAN) and Wide Area Network (WAN) environment. They should also know how to optimize EIGRP with the use of route summarization.

Objectives

Upon completing this lesson, you will be able to:

- Describe how EIGRP operates
- Describe how EIGRP builds and maintains neighbor relationships
- Configure EIGRP
- Use summary addressing to limit the scope of EIGRP queries
- Control EIGRP split horizon issues

Learner Skills and Knowledge

To fully benefit from this lesson, you must have these prerequisite skills and knowledge:

- Passed the CCIE written qualification exam
- Completed the Building Scalable Cisco Internetworks (BSCI) course or have the equivalent knowledge

Outline

This lesson includes these sections:

- Overview
- What is EIGRP?
- Configuring EIGRP
- EIGRP Route Summarization
- Load Balancing with EIGRP
- EIGRP Split Horizon
- Verifying EIGRP Operation
- Summary
- Lesson Assessment (Quiz)

What is EIGRP?

Enhanced Interior Gateway Routing Protocol (EIGRP) is a classless routing protocol that directly interfaces to Internet Protocol (IP) as protocol 88.

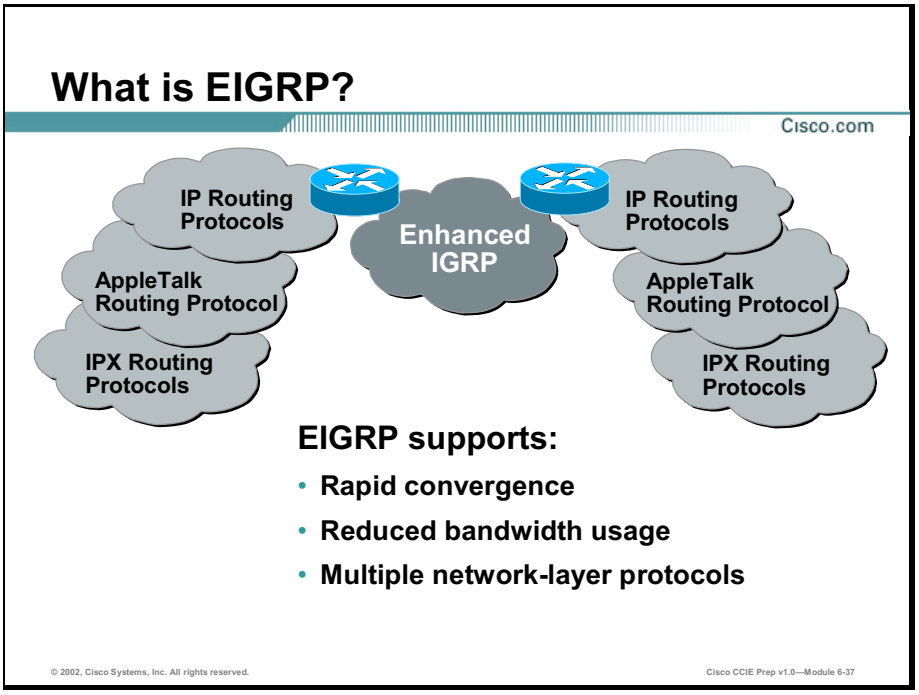

EIGRP uses the multicast address of 224.0.0.10 for 'hellos' and routing updates instead of an all hosts broadcast like Routing Information Protocol (RIP) uses. EIGRP also employs a system of hello and hold timers to maintain neighbors. Aside from the initial routing update, partial routing updates are sent only when network topology changes occur. The updates are also "bounded", which means updates are sent only to pertinent routers. Like IGRP, EIGRP uses a composite metric to calculate the best path to a destination. The sections that follow take a closer look at how EIGRP makes use of metrics, neighbors, reliable transport, and Diffusing Update Algorithm (DUAL) in its operation. Some EIGRP features are as follows:

- **Support for Variable Length Subnet Mask (VLSM):** EIGRP is a classless routing protocol, and carries the subnet mask of the route in its update.

- **Rapid convergence:** By using the concept of feasible successors, defined by DUAL, EIGRP is able to pre-select the next best path to a destination. This allows for very fast convergence upon a link failure.

- **Low Central Processing Unit (CPU) utilization:** Under normal operation only 'hellos' and partial updates are sent across a link. Routing updates are not flooded and processed only periodically.

- **Incremental updates:** EIGRP does not send full routing updates, it only sends information about changed routes.

- **Scalable:** Through the use of VLSM and a complex composite metric, EIGRP networks can scale dramatically in size.

- **Easy configuration:** EIGRP supports hierarchical network design, but it does not require the strict configuration guidelines, such as the ones needed for Open Shortest Path First (OSPF).

- **Automatic route summarization:** EIGRP will perform automatic summarization on major bit boundaries.

- **Message Digest Version 5 (MD5) route authentication:** (as of Cisco Internetwork Operating System (IOS) Software Release 11.3) EIGRP can be configured to perform MD5 password authentication on route updates.

Configuring EIGRP

This section discusses the configuration of EIGRP.

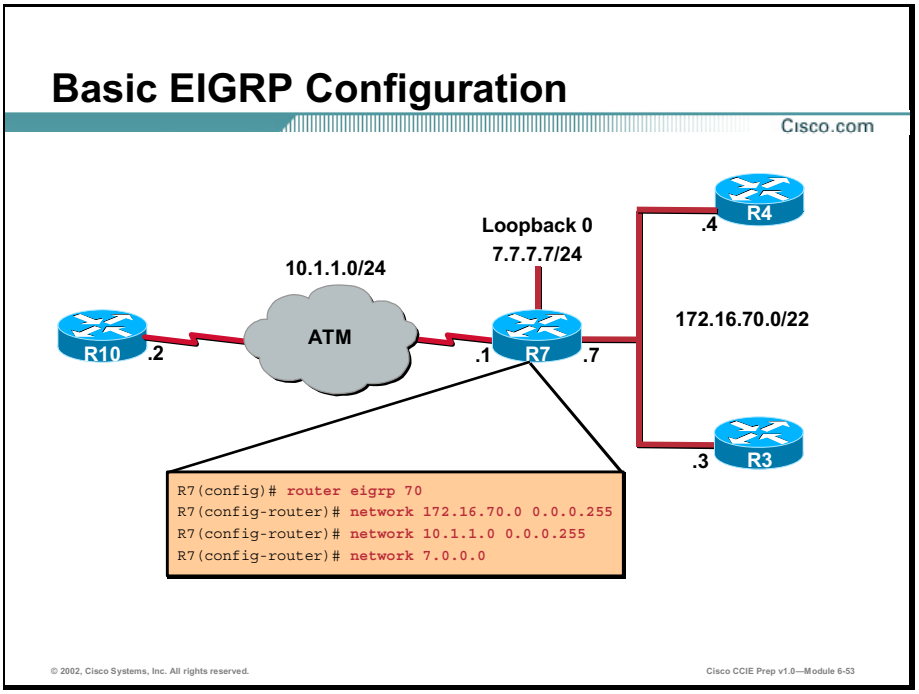

Configuring EIGRP calls for the definition of an Autonomous System (AS). By definition, an AS is a set of routers under a single administrative technical authority. Like IGRP, EIGRP uses the concept of autonomous systems to separate routing processes. Having a registered AS when configuring EIGRP is not required, nor does EIGRP use the AS for routing decision.

The following two-step process can be used to configure EIGRP.

Step 1 Enable EIGRP and define an autonomous system on the router. This is accomplished with the **router eigrp** *autonomous_system_id* global command.

Step 2 Add the networks you want to be included in the EIGRP routing process. This is accomplished with the **network** *a.b.c.d* from the config-router# mode. When you enter the network statements, it is only necessary to enter the major class boundary. In Cisco IOS Software Release 12.0 and later, the **network** command adds an additional wildcard mask, much like OSPF. This is an inverse bit mask, for example to enable EIGRP on network 172.16.70.0 only; the syntax would be **network 172.16.70.0 0.0.0.255.**

Note	EIGRP converts a subnet mask to a wildcard mask if you make an error.

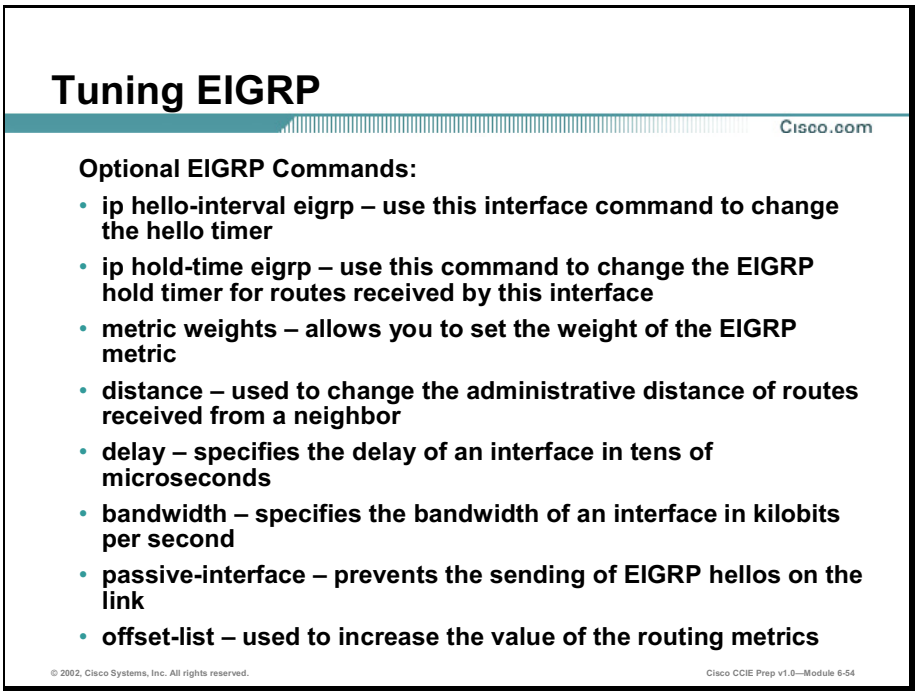

The following is a list of parameters adjustable for EIGRP.

- router(config-if)# **ip hello-interval eigrp** *as_number interval_in_seconds* Use this interface command to change the hello timer for EIGRP. The default value of this command is interface dependant. By default, hello packets are sent every 5 seconds. The exception to this is low-speed, Nonbroadcast Multi-Access media (NBMA), where it is 60 seconds. Low-speed is defined as rates of T1 (1.544 Megabits per second (Mbps)) or slower. All neighbors residing on a network should have equal hello timers.

- router(config-if)# **ip hold-time eigrp** *as_number holdown_timer_in_seconds* Use this command to change the EIGRP hold timer for routes received by this interface. The timer has a default vault of 180 seconds for low-speed NBMA networks, and 15 seconds for all other networks. All neighbors residing on a network should have an equal hold timer.

The following subsets of commands are used to influence routing decisions made by EIGRP. Individual metrics may be modified as well as the administrative distance of the EIGRP. Whenever influencing a specific link's metric, use the **delay** command over the **bandwidth** command. Both may be used, but remember that OSPF will also be affected by **bandwidth,** while **delay** will affect only IGRP and EIGRP.

- router(config-router)# **metric weights 0** *k1 k2 k3 k4* k5 This command will allow you to set the weight of the EIGRP metric, in terms of bandwidth, load, delay, and reliability. Change these values with extreme caution, as EIGRP will not form neighbors with mismatched K values.

- router(config-router)# **distance** [*1-255*] *adjacent_neighbors_ip_address wildcard_mask* [*access_list_0-99]* Use this command to change the administrative distance of routes

received from a neighbor. If the IP address and wildcard_mask are omitted, all routes for that protocol will be set to the distance value.

- router(config-if)# **delay** [*[ms] 1-4214748364*] Specifies the delay of an interface in tens of microseconds. This command is used only by routing protocols, and does not affect traffic on the link.

- router(config-if)# **bandwidth** [*bandwidth_kbps 1-4214748364*] Specifies the bandwidth of an interface in kilobits per second. This command is used only by routing protocols, and does not affect traffic on the link. The bandwidth parameter should be set on all interfaces to give EIGRP an accurate view of the network.

- router(config-router)# **passive-interface** *interface_name* Prevents the sending of EIGRP hellos on the link. This command operates differently on EIGRP then IGRP. Because hellos are suppressed, no neighbors will be formed; therefore, no routing updates will be sent or received.

- router(config-router)# **offset-list** [*access_list_0-99* {**in** | **out**} **offset** [*metric_offset_1-214748364*] [*interface*] Use this to increase the value of the routing metrics. The metric offset cannot exceed 214748364. The offset list is applied in the same way as it is in RIP, using the EIGRP metric.

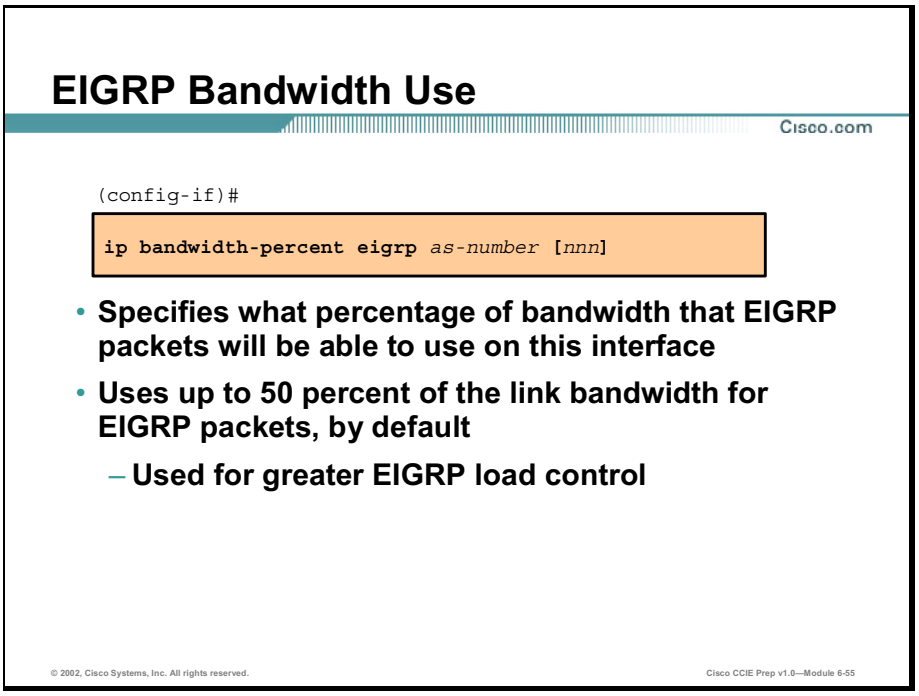

By default, EIGRP will use up to 50 percent of the bandwidth of an interface or subinterface, as set with the **bandwidth** parameter. You can change this percentage on a per-interface basis by using the following interface command:

```
router(config-if)# ip bandwidth-percent eigrp [as-number] [nnn]
```

In this command, *nnn* is the percentage of the configured bandwidth that EIGRP is allowed to use. Note that this percentage can be set to greater than 100. This capability is useful if the bandwidth is configured artificially low for routing policy reasons. For example:

```
interface serial0
bandwidth 20
ip bandwidth-percent eigrp 1 200
```

This configuration would allow EIGRP to use 40 kilobits per second (kbps) (200 percent of the configured bandwidth) on the interface. It is essential to make sure that the line is provisioned to handle the configured capacity.

EIGRP Route Summarization

Summarization provides two powerful enhancements to EIGRP. First by lowering the number of routes in the route table, it lessens the amount and size of the EIGRP advertisements. Secondly, and more importantly, it can limit the EIGRP query range.

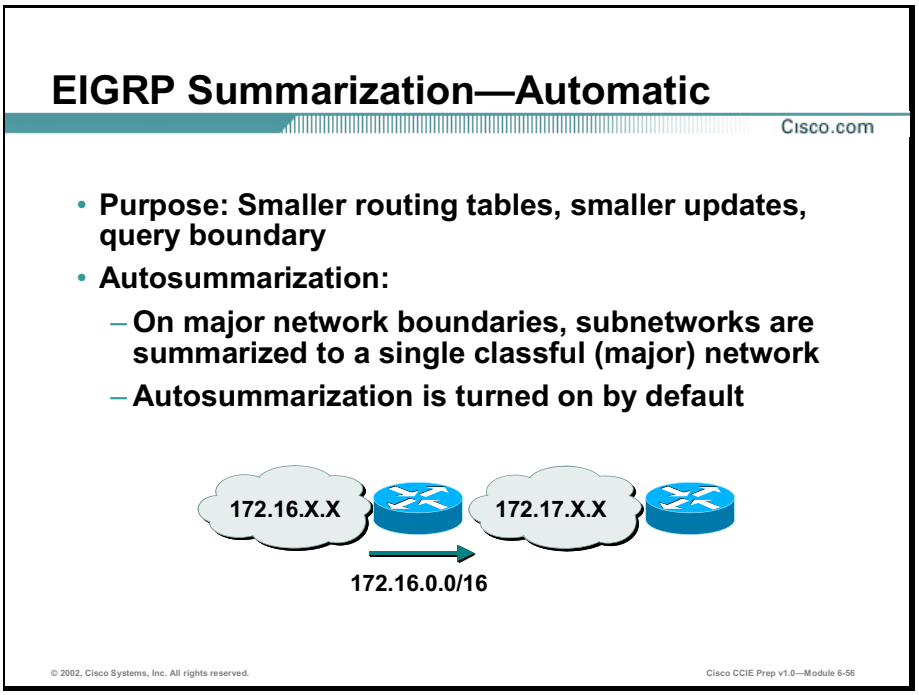

By default, EIGRP performs auto summarization in two situations:

- Auto summarization will occur at the major class boundary during redistribution from EIGRP into a classful routing protocol such as IGRP or RIP. This type of summarization cannot be disabled.
- Auto summarization will occur at the major class boundary when the route is advertised out an interface that is on a different major class boundary. This summarization can be disabled with the **no auto-summary** command in EIGRP router configuration mode.
- EIGRP will not automatically summarize EIGRP external routes.
- EIGRP routes that are summarized will have an administrative distance of 90.

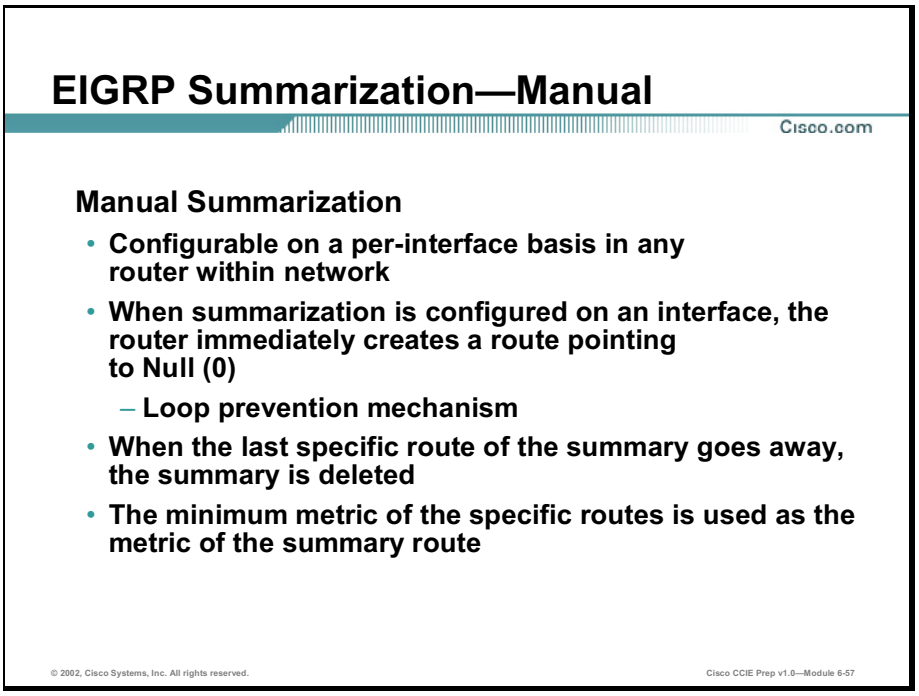

EIGRP manual summarization is critical to large EIGRP networks. It limits the EIGRP query and can significantly reduce the size of the routing table. There are essentially two ways to deploy manual summarization:

Advertise a summary address or aggregate address with the following interface command:

ip summary-address eigrp *as_number summary_address address_mask*

Advertise a default summary route, with the following interface command:

ip summary-address eigrp *as_number* **0.0.0.0 0.0.0.0**.

This command will cause only the default route to be advertised and all other routing updates will be suppressed.

Note	In Cisco IOS 12.0(4)T, an administrative distance can be added to the summary address to alter the default admin distance of 90.

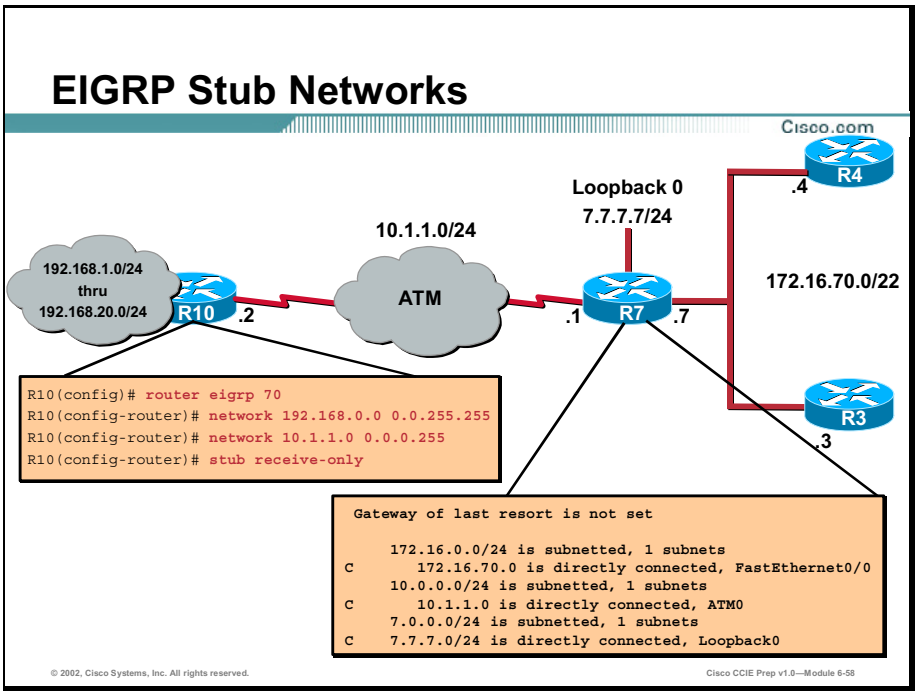

In Cisco IOS Software Release 12.0(7)T, Cisco introduced EIGRP stub routing to further control stability and reduce resource utilization. This feature was fully integrated into Release 12.0(15)S. EIGRP stub routing functions like that of an OSPF stub area. The stub router will have one exit path from the routing domain and forward all traffic to a central or distribution router. Another way to say this is that the stub network cannot be a transit router for EIGRP, and it can have only one EIGRP neighbor.

When configuring EIGRP stub routing, only the remote or the spoke router needs to be configured as a stub. This router will respond to queries for summaries, connected routes, redistributed static routes, external routes, and internal routes with the message “inaccessible”. This process greatly reduces the overhead associated with responding to queries by the remote routers. The stub router will also send special peer information to its neighbor informing its neighbor that it is a stub router.

To configure an EIGRP stub routing, use the following router command under EIGRP.

```
router(config-router)# eigrp stub [receive-only | connected | static | summary]
```

The options are described as follows:

- **Receive-only**: Causes the router to not send any routes.
- **Connected**: The router advertises all connected routes to the single neighbor. No redistribution is necessary.
- **Static**: The router advertises all static routes to a single neighbor. The static routes still need to be redistributed into EIGRP to be advertised.
- **Summary**: The router advertises summary routes.

To verify that the router is configured as an EIGRP stub router, use the **show ip eigrp neighbor detail** command. The last line of the output will show if stub routing is enabled, and what the stub router can advertise. The **show eigrp packet stub** will show debug information about the stub status of the peer routers.

Load Balancing with EIGRP

EIGRP supports equal cost and unequal cost load balancing over a maximum of six paths.

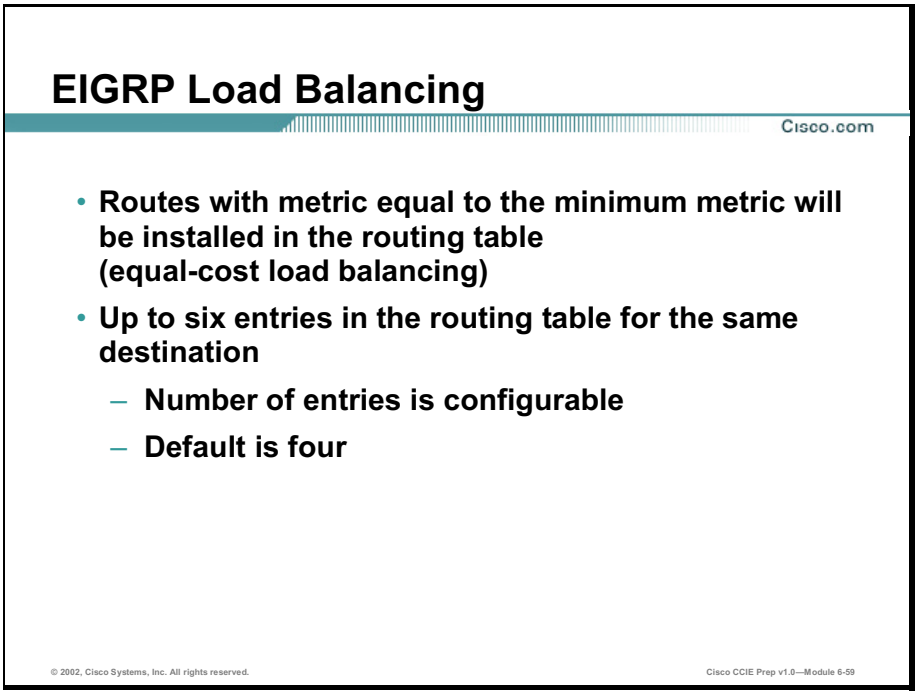

By default, EIGRP will load-share over four equal cost paths. For load-sharing to happen, the load-sharing routes must show up in the IP forwarding table, or with the **show ip route** command. Only when a route shows up in the forwarding table with multiple paths to it, will load sharing occur. Use the **bandwidth** interface command on serial links to ensure EIGRP has a consistent perspective of the metrics of the network. This may also aid in making the route show up in the IP forwarding table.

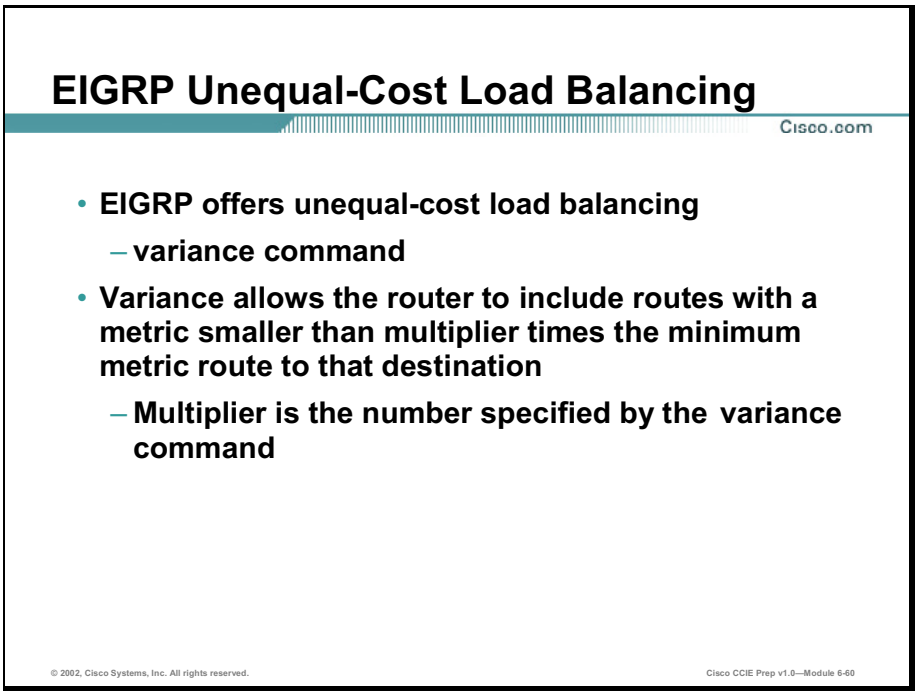

EIGRP also has the capability to use unequal-cost load balancing in the same manner as IGRP. The router uses variance as a multiplier in choosing the upper boundary of path with the greatest metric.

Configuring EIGRP unequal-cost load balancing is a three-step process:

Step 1 Configure the bandwidth on both sides of all the interfaces involved in the load-sharing group. Use the **bandwidth** *xx_kbps* command to accomplish this.

Step 2 Define the lowest cost metric and the highest cost metric. From these values, compute the variance multiplier and add it to the EIGRP routing process. The composite metric EIGRP is using can be viewed with the **show ip eigrp topology** command, as discussed in previous sections.

Step 3 (Optional) Set the *maximum-paths* or the *traffic-share* variables.

The following example takes you through the calculation of a fictional variance. EIGRP has a route and the metric of that route is 100. The router also has two more routes to that same destination, and the metric for those routes is 200 and 300. To allow EIGRP to use all three paths in sharing data, set the variance to 3.

(3 * 100) = 300.

Another way to view it is the (lowest_metric) = largest metric of path to load share over, in this case 300. To properly set the variance in a real network use the following formula:

```
Variance = 1 + [[[metric of highest cost route] / [metric of the lowest cost
route]] rounded up to the nearest 1s decimal place.
```

The metric of the lowest cost and highest cost routes can be discovered with the **show ip eigrp topology** command. Be sure to change variance and any other variables, such as bandwidth, on both ends of the link. The bandwidth should be set on all serial links. The following are the syntax for the commands used in configuring load balancing:


```
router(config-router)# variance [metric_multiplier 1-128]
router(config-router)# maximum-paths [1-6]
router(config-router)# traffic-share [balanced | min across-interface]
router(config-if)# bandwidth xx kbps
```

The **variance** command defines the metric multiplier of which routes to use in unequal-cost load balancing. The default variance is 1, which is equal cost load balancing.

By default the router will load balance across four equal cost paths. To modify this number use the **maximum-paths** command. The maximum setting for this command is six equal cost paths. The minimum setting of one disables load balancing. EIGRP can perform unequal cost load balancing in addition to equal cost load balancing.

The multiple paths that make up a single hop transport to a common destination are called a load-sharing group. The default value is 4.

With the **traffic-share** command, if there are multiple minimum-cost paths and **traffic-share-min** is configured, EIGRP will use equal-cost load balancing. By default, the command is set to **balanced**, where traffic will be distributed proportionally to the ratio of the metrics. For example, if variance is set to 3, and traffic-share is set to balanced, then the best route will transport traffic three times that of the worst route.

For a route to be included in unequal-cost load sharing, three other conditions must be met.

- The maximum-paths limit must not be exceeded as a result of adding this route to the load-sharing group.
- The downstream router must be metrically closer to the destination.
- The metric of the lowest-cost route, multiplied by the variance, must be greater than the metric of the route to be added to the load-sharing group.

EIGRP Split Horizon

EIGRP actually runs its own version of split horizon for all of the protocols that it supports: IP, Internetwork Packet Exchange (IPX), and Appletalk.

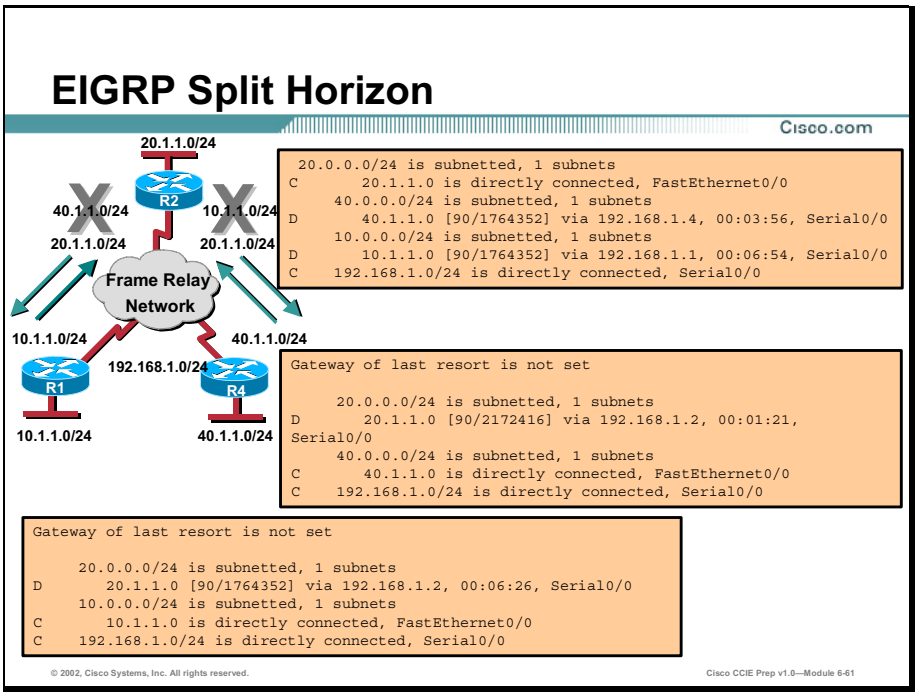

Frame Relay Technologies and split horizon are a routing technique used with classful routing protocols in which routing updates are prevented from being advertised out the same interface from which they were learned. Split horizon issues are most prevalent in NBMA hub and spoke networks. To remedy this, split horizon is disabled by default on physical interfaces and point-to-multipoint subinterfaces when they are configured for Frame Relay encapsulation.

This default setting alleviates reachability issues, but can cause other problems that you need to be aware of, such as routing loops. This default behavior applies to RIP and IGRP, however EIGRP runs its own version of split horizon that must be explicitly disabled for both IP and IPX in a NBMA hub and spoke environment. An alternative to disabling split horizon is the use of point-to-point subinterfaces.

As shown in the example, R2 receives routing updates from R4 and R1, but because of split horizon, R2 does not advertise the 40.0.0.0/24 network to R1 and the 10.0.0.0/24 network to R4. Therefore, the remote sites do not have full reachability.

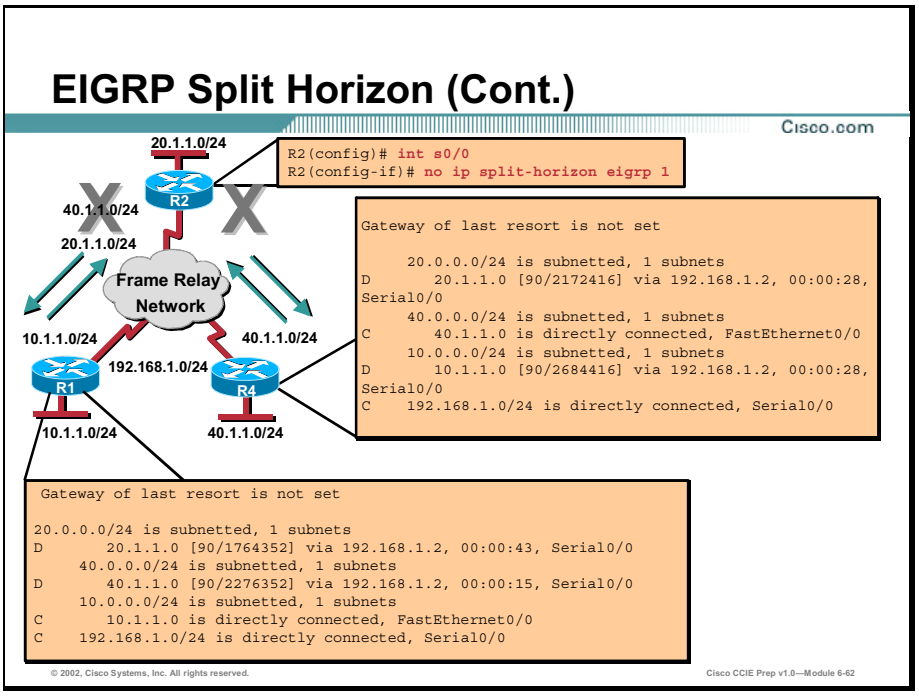

To allow full reachability between the remote sites, you must disable split horizon on R2. This is done with the **no ip split-horizon eigrp** *autonomous system* command. This is an interface configuration command that must be placed on the interface for which you wish to disable EIGRP split horizon.

Notice in the example that the spoke routers now have full routing tables and therefore can reach all remote sites.

Verifying EIGRP Operation

This section describes how to verify EIGRP operation.

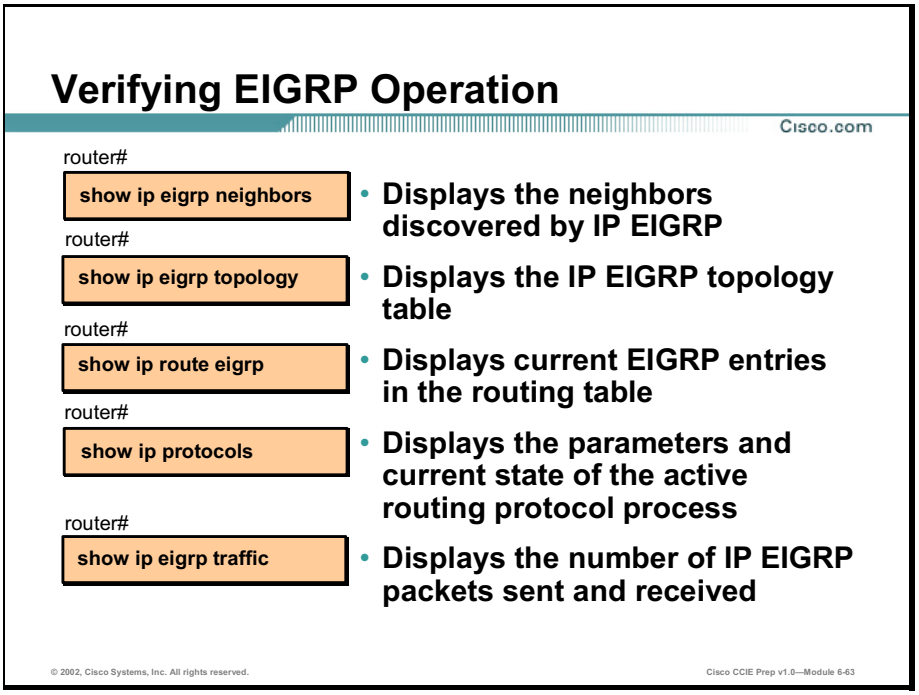

The **show** commands can be used to verify EIGRP operation.

Table 5-1: show Commands

Command	Description
`show ip eigrp neighbors`	Displays neighbors discovered by EIGRP
`show ip eigrp topology`	Displays the EIGRP topology table. This command shows the topology table, the active or passive state of routes, the number of successors, and the FD to the destination.
`show ip route eigrp`	Displays the current EIGRP entries in the routing table
`show ip protocols`	Displays the parameters and current state of the active routing protocol process. This command shows the EIGRP AS number. It also displays filtering and redistribution numbers, as well as neighbors and distance information.
`show ip eigrp traffic`	Displays the number of EIGRP packets sent and received. This command displays statistics on hello, updates, queries, replies, and acknowledgments.

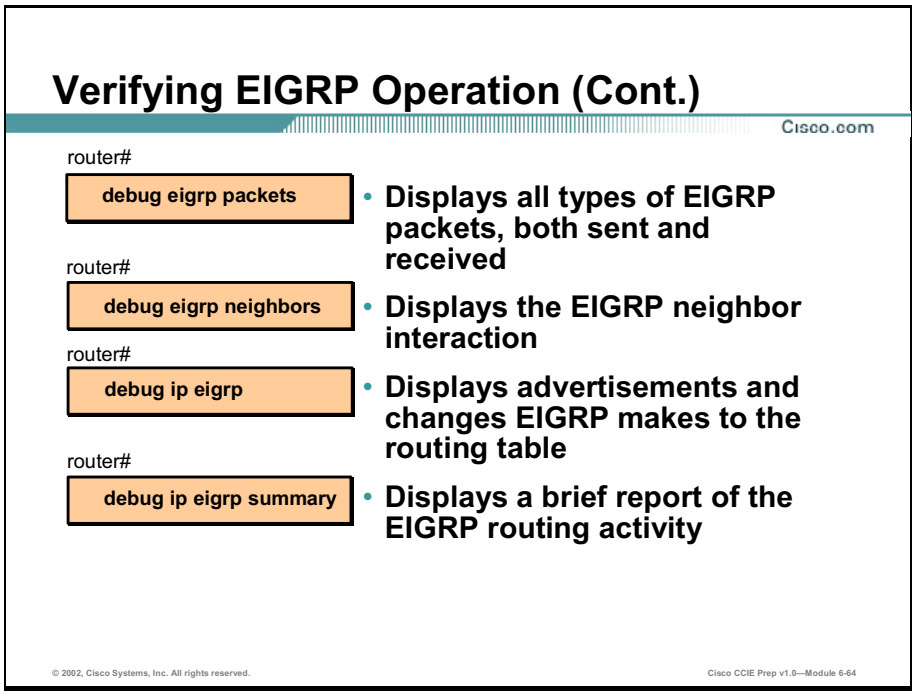

These **debug** commands can be used to verify EIGRP operation.

Table 5-2: debug Commands

Command	Description
`debug eigrp packets`	Displays the types of EIGRP packets sent and received. A maximum of 11 packet types can be selected for individual or group display.
`debug eigrp neighbors`	Displays the neighbors discovered by EIGRP and the contents of the hello packets
`debug ip eigrp`	Displays EIGRP packets that are sent and received
`debug ip eigrp summary`	Displays a summarized version of EIGRP activity. It also displays filtering and redistribution numbers, as well as neighbors and distance information.

show ip eigrp neighbors Command

This can be one of the most useful commands when verifying the operational status of EIGRP. The **show ip eigrp neighbors** command will show the status of all EIGRP neighbors. The neighbor should be "up" for as long as EIGRP has been running on the link. EIGRP will form a neighbor with all routers on the same subnet, and in the same AS. EIGRP will not form a neighbor with mismatched *k* values; but a neighbor can be formed with mismatched hellos and dead timers. A neighbor with a short uptime is a clear indication of a problem. Another important field is the **queue count**. This field indicates the number of packets waiting to be transmitted to that neighbor. This value should be 0 or a number under 20. Consistent Q values in the range of 60 or greater are considered high. A high Smooth Round Trip Timer (SRTT) number can mean the packet is experiencing some type of delay on the link.

```
R1# show ip eigrp neighbors
IP-EIGRP neighbors for process 2001
H   Address                 Interface   Hold Uptime   SRTT   RTO  Q  Seq
                                        (sec)         (ms)       Cnt Num
1   172.16.1.5              Se0.1        136 05:48:23   36  1302  0  15
0   172.16.1.6              Se0.1        131 05:48:24   40  1302  0  17
R1#
```

- **Handle (H):** A Cisco IOS internal number used to identify a neighbor. Do not confuse this with hop count.
- **Neighbor Address:** This is the adjacent neighbor's IP address. A neighbor should be formed between every router on that subnet running EIGRP in a common AS.
- **Interface:** The interface that is reporting the neighbor.
- **HoldTime:** This is the amount of time, which counts down, that EIGRP waits for a 'hello' before tearing down the neighbor.
- **Uptime:** States how long the neighbor has been up. This number should be up for as long as the link has been up.
- **SRRT:** The number of milliseconds it takes for an EIGRP packet to be sent to this neighbor, and for the local router to receive an acknowledgement, hence, a round trip timer. If this number equals zero, a packet has never made a successful round trip.
- **Retransmission TimeOut (RTO):** The amount of time, in milliseconds, that the EIGRP waits before re-transmitting a packet from the retransmission queue to a neighbor.
- **Queue count (Q):** The number of packets waiting in the queue to be sent out to this neighbor. This value should be 0 or a very low number. A high queue count indicates that data is having trouble getting through.
- **Sequence Number (Seq-Num):** Sequence number of the last update, query, or reply that was received from this neighbor. If this number equals zero, it indicates that no reliable packets have ever been received from the neighbor, another clear indication of a problem.

show ip eigrp topology Command

This command lists the EIGRP topology table discussed earlier. The table lists all routes that EIGRP is aware of and whether EIGRP is actively processing information on that route. Under most normal conditions, the routes should all be in a passive state, no EIGRP processes are running for that route. If the routes are active, this could indicate the Stuck In Active (SIA) state, which will be discussed in more detail in an upcoming section. The **show ip eigrp topology** command can also be extended to show information about an individual route or subnet. This information will include all relevant information about the route, including all of

its metrics and successors, as well as how the route was learned. This example illustrates the use of **show ip eigrp topology** followed by the extended version of the command.

```
R1# show ip eigrp topology
IP-EIGRP Topology Table for process 2001

Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - Reply status

P 172.16.5.0/24, 1 successors, FD is 23394560
         via 172.16.1.5 (23394560/281600), Serial0.1
P 172.16.6.0/24, 1 successors, FD is 23394560
         via 172.16.1.6 (23394560/281600), Serial0.1
P 172.16.1.0/24, 1 successors, FD is 23368960
         via Connected, Serial0.1
P 172.16.2.0/24, 1 successors, FD is 281600
         via Connected, Ethernet1
R1#

R1# show ip eigrp topology 2001 172.16.5.0 255.255.255.0
IP-EIGRP topology entry for 172.16.5.0/24
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 23394560
  Routing Descriptor Blocks:
  172.16.1.5 (Serial0.1), from 172.16.1.5, Send flag is 0x0
      Composite metric is (23394560/281600), Route is Internal
      Vector metric:
        Minimum bandwidth is 112 Kbit
        Total delay is 21000 microseconds
        Reliability is 254/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 1
R1#
```

The fields to note in this output are as follows:

- **P -** Passive, no EIGRP computation is being performed. This is the ideal state.
- **A -** Active, EIGRP computations are “actively” being performed for this destination. Routes constantly appearing in an active state, indicates a neighbor or query problem. Both are symptoms of the SIA problem.
- **U -** Update, an update packet was sent to this destination.

- **Q -** Query, a query packet was sent to this destination.
- **R -** Reply, a reply packet was sent to this destination.
- **Route information -** IP address of the route or network, its subnet mask, and the successor, or next hop to that network, or the feasible successor.
- **FD -** Feasible distance to the destination network.
- **Send Flag -** The type of packets that need to be sent for the entry are indicated by the send flag.

  0x0 – If there are packets tha need to be sent in relation to this entry, this indicates the type of packet.

  0x1 - The router has received a query for this network and needs to send a unicast reply.

  0x2 - The route is active and a multicast query should be sent.

  0x3 - The route has changed and a multicast update should be sent.

Summary

This section summarizes the key points discussed in this lesson.

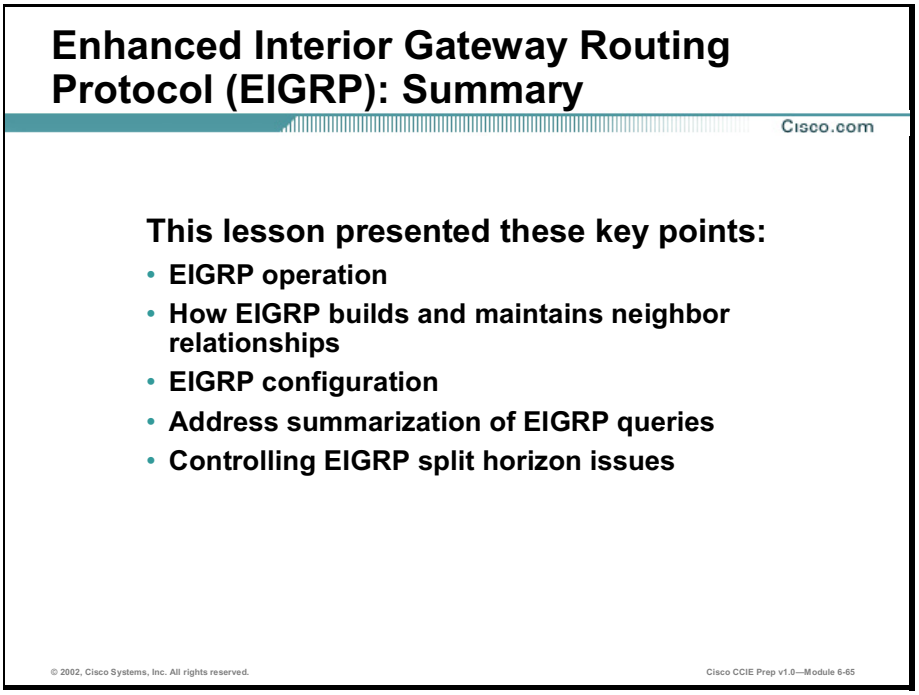

Next Steps

After completing this lesson, go to:

- Link-State Routing Protocols

References

For additional information, refer to these resources:

- *Routing TCP/IP Volume I* by Jeff Doyle

Lesson Assessment (Quiz)

Q1) True or False: If EIGRP passive interface is enabled, EIGRP will still receive routes, but it will not advertise any.

A) True

B) False

Q2) True or False: EIGRP is not susceptible to split horizon issues.

A) True

B) False

Q3) True or False: EIGRP will not establish a relationship with a neighbor with mismatched timers.

A) True

B) False

Q4) By default, EIGRP uses the following metrics on which to base its routing decisions.

A) MTU, Bandwidth, Load

B) MTU, Delay

C) MTU, Bandwidth, Load, Reliability, Delay

D) Bandwidth, Delay

6

Link-State Routing Protocols

Overview

This module describes the configuration of a common link-state protocol, Open Shortest Path First (OSPF). OSPF can be configured in a Single Area or in Multiple Areas. This module describes the configuration of both.

Upon completing this module, you will be able to:

- Configure OSPF in a multi-area environment
- Configure advanced OSPF features, such as neighbor authentication, demand circuits, and virtual links
- Verify the operation of OSPF using various **show** and **debug** commands

Outline

The module contains these lessons:

- Configuring OSPF in a Single Area
- Multi-Area OSPF Environments
- Advanced OSPF Features
- Troubleshooting OSPF

Configuring OSPF in a Single Area

Overview

A single Open Shortest Path First (OSPF) area can encompass enterprise-size companies and may involve numerous network topology considerations. This lesson examines configuring OSPF in a single area and also describes the detailed configuration required by OSPF for each of the common Wide Area Network (WAN) topologies used by companies today.

Importance

OSPF is the core Interior Gateway Protocol (IGP) used in the Cisco Certified Internetworking Expert (CCIE) lab.

Objectives

Upon completing this lesson, you will be able to describe:

- OSPF operation and configuration in a broadcast multi-access environment
- OSPF operation and configuration in a point-to-point topology
- OSPF operation and configuration in a Non-Broadcast Multi-Access (NBMA) environment

Learner Skills and Knowledge

To fully benefit from this lesson, you must have these prerequisite skills and knowledge:

- Passed the CCIE written qualification exam
- Completed the Building Scalable Cisco Internetworks (BSCI) course or have the equivalent knowledge

Outline

This lesson includes these sections:

- Overview
- OSPF Configuration in a Broadcast Multi-Access Topology
- Controlling the Designated Router/Backup Designated Router (DR/BDR) Election
- OSPF Operation in an NBMA Topology
- Summary
- Lesson Assessment (Quiz)

OSPF Configuration in a Broadcast Multi-Access Topology

This section discusses Open Shortest Path First (OSPF) operation and configuration in a broadcast multi-access environment such as Ethernet or Token Ring.

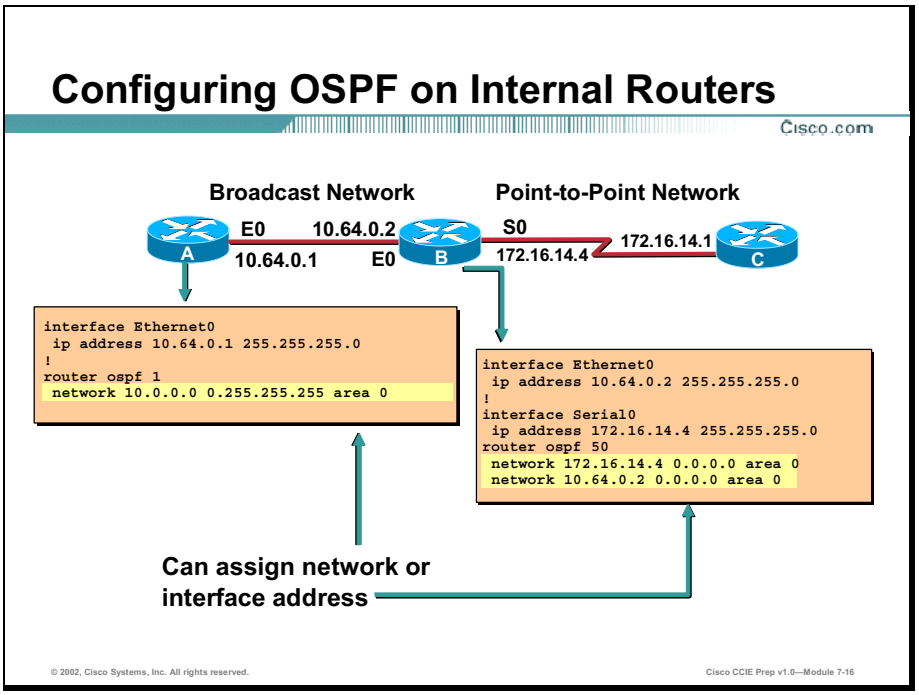

As with other routing protocols, enabling OSPF requires that you create an OSPF routing process and specify the networks to be associated with the routing process. However, OSPF requires the use of an inverse mask with the **network** command to control exactly which interfaces on the router participate in OSPF. Also, the interfaces specified with the **network** command must be assigned to a particular area with the **area** parameter.

Table 6-1: OSPF Routing Process

Steps	Command	Purpose
Step 1:	**router ospf** *<process-id>*	Enables OSPF routing on the router. The process ID is an internally used number to identify the OSPF processes running on the router. The process-id does not need to match process-ids on other routers to share routing information. Running multiple OSPF processes on the same router is not recommended because it creates multiple database instances that add extra overhead.
Step 2:	**network** *address wildcard-mask* **area** *area-id*	Defines the interfaces on which OSPF will run and the area ID for those interfaces. *address* - Can be the network address, subnet, or the actual Internet Protocol (IP) address of the interface. This parameter instructs the router on which interfaces to send and listen for LSAs and what networks to advertise. *wildcard-mask* - An inverse mask used to determine the range of interfaces to run OSPF on. The mask has wildcard bits, where 0 is a match and 1 is the "don't care" bit. For example, 172.16.0.0 0.0.255.255 indicates a match in the first two bytes of 172.16. The router will enable OSPF on all interfaces that fall within the 172.16.x.x range. If specifying an actual interface address, use the mask 0.0.0.0 to match all four bytes of the address. An address and wildcard-mask combination of 0.0.0.0 255.255.255.255 will match all interfaces on the router. **area** - Specifies the area that the interfaces will be assigned to. This parameter can be specified in decimal (0-65,535) or dotted decimal (A.B.C.D, similar to an IP address) format.

Controlling the DR/BDR Election

This section covers how to control the DR/BDR election.

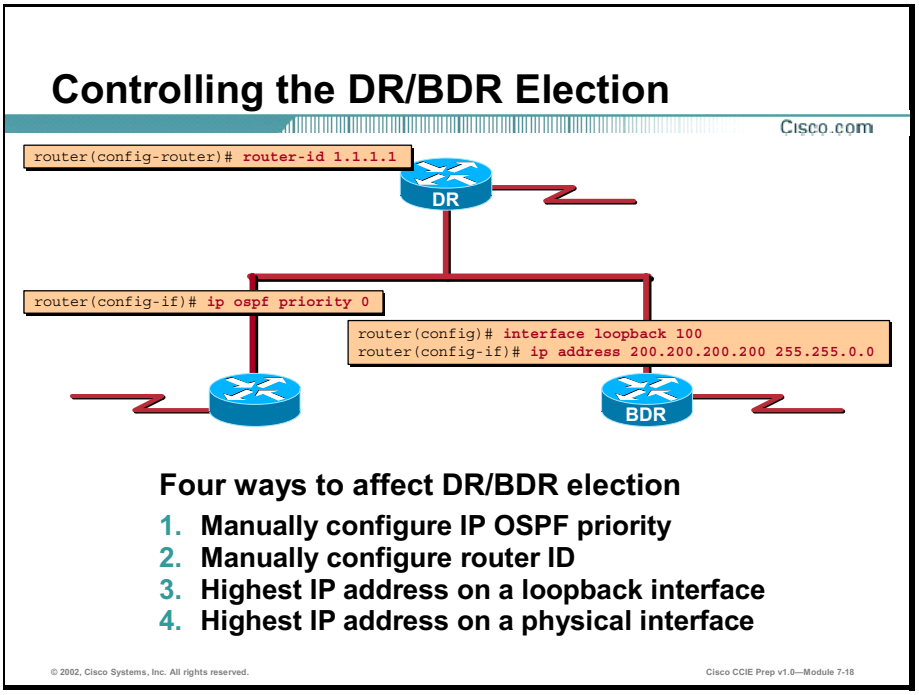

OSPF uses the highest Internet Protocol (IP) address configured on an interface as its router ID. If the interface associated with this IP address is ever brought down, or if the address is removed, the OSPF process must recalculate a new router ID and resend all of its routing information back out.

If a loopback interface is configured with an IP address, the Cisco IOS software will use this IP address as its router ID, even if other physical interfaces have higher IP addresses. Due to the fact that loopback interfaces never go down, greater stability in the routing table is achieved using a loopback interface as the router ID.

OSPF automatically prefers a loopback interface over a physical interface of any kind, and it chooses the highest IP address among all loopback interfaces if multiple loopbacks exist. If no loopback interfaces are present, the highest IP address on a physical interface is chosen as the router ID. You cannot tell OSPF to use any particular interface as the router ID; however, you can manually set the router ID with the **router-id** command.

The best way to control the DR/BDR election is to manually change the priority on the interface. The default priority on an interface is 1. Higher values have higher priority. This value can be set to any number from 0 – 255. A setting of 0 prevents the router from participating in the DR/BDR election.

Table 6-2: DR/BDR Election Commands

Command	Description
interface loopback 0	Creates a virtual loopback interface on the router
ip address ip-address mask	Assigns an IP address to this interface
router-id A.B.C.D	Allows you to manually set the router ID on a router. The router ID does not have to be an IP address that exists on the router, but must be in dotted decimal format, similar to an IP address. Performed from router configuration mode
ip ospf priority <0-255>	Interface configuration command that controls the likelihood of the router becoming the DR or BDR for a network segment

Note Once the OSPF process is started on a router and neighbor adjacencies have been formed, the router ID will not change. For example, if an IP address is added to the router that is higher than the current router ID, the router ID will not change. However, the router ID will change to the new IP address when the OSPF process is restarted or the router is reloaded. This will invalidate items such as virtual-links that were configured with the old router ID. Virtual-links will be covered in the next lesson.

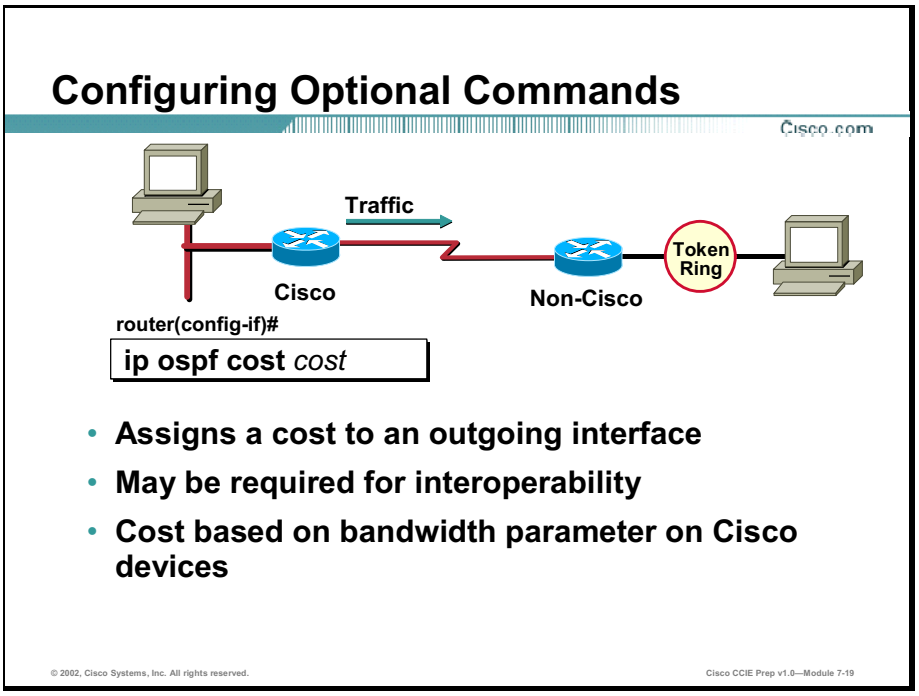

Cisco's OSPF default cost assignment is based on the bandwidth of the link. Other vendors might use a different mechanism to assign OSPF cost to a link, so you may have to manually set the cost associated with a link in some scenarios. OSPF requires that all interfaces connected to a link agree on the link's cost. By default, Cisco routers calculate the cost of a link using the following formula.

- Reference Bandwidth / Bandwidth

The default reference bandwidth is Fast Ethernet (100 Mbps). Therefore, the formula can also be written as:

- 100,000,000 / Bandwidth

Using this formula, here are some examples that yield the default costs:

- **56-Kilobits per second (Kbps) serial link:** Default cost is 1785
- **T1 (1.544-Megabits per second (Mbps) serial link):** Default cost is 64
- **Ethernet:** Default cost is 10
- **16-Mbps Token Ring:** Default cost is 6

Table 6-3: <ip ospf cost> Commands

Command	Description
ip ospf cost <1-65,535>	Interface configuration command that assigns a value from 1 to 65535 that indicates the cost of the interface. The cost of a route in OSPF is the sum of the costs of all outgoing interfaces to that destination.

Note On serial lines, the default bandwidth is 1.544 Mbps. If the link's speed is actually slower than that, use the **bandwidth** command to specify the real link speed. This will ensure accurate metrics in routing protocols, such as OSPF, Interior Gateway Routing Protocol (IGRP), and Enhanced IGRP (EIGRP).

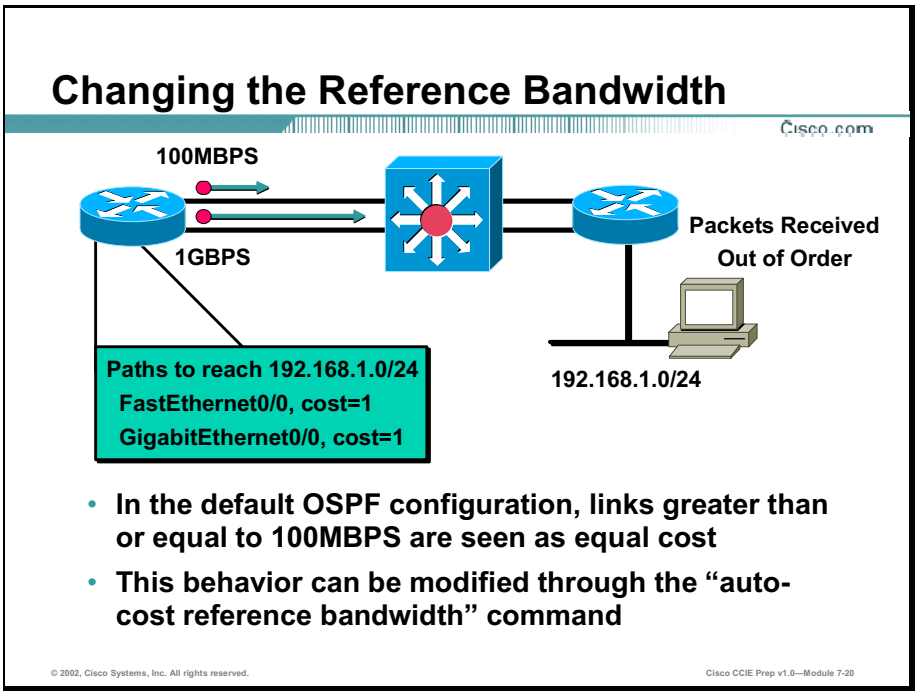

When OSPF was developed in the early 1990s, the designers seemed to believe that the fastest media that would ever be available was 100 Mbps. At least, that is how they wrote the formula that OSPF uses to calculate costs. This becomes a problem when you add some newer technologies such as Asynchronous Transfer Mode (ATM) and Gigabit Ethernet into your OSPF network. Using the default reference bandwidth of 100,000,000, OSPF will see all of these link types including Fast Ethernet as a cost of 1. For example, this becomes a problem when a certain destination is accessible over both Gigabit Ethernet and Fast Ethernet. OSPF will actually try and load balance across these two links. If you are running a media type that is faster than 100 Mbps you should actually change the reference bandwidth that OSPF uses in its formula for calculating costs. Since Gigabit Ethernet is now widely available it is suggested that you change the reference bandwidth to reflect Gigabit Ethernet speeds, which correlates to 1,000,000,000.

Table 6-4: < auto-cost > Commands

Command	Description
auto-cost reference-bandwidth <ref-bw>	Reference bandwidth, in Mbps. The range is 1 to 4294967; the default is 100. Performed from router configuration mode.

Note Any change using this command must be done on all routers in the autonomous system so that they are all using the same formula to calculate cost. The value set by the **ip ospf cost** command still overrides the cost resulting from the **auto-cost reference-bandwidth** command.

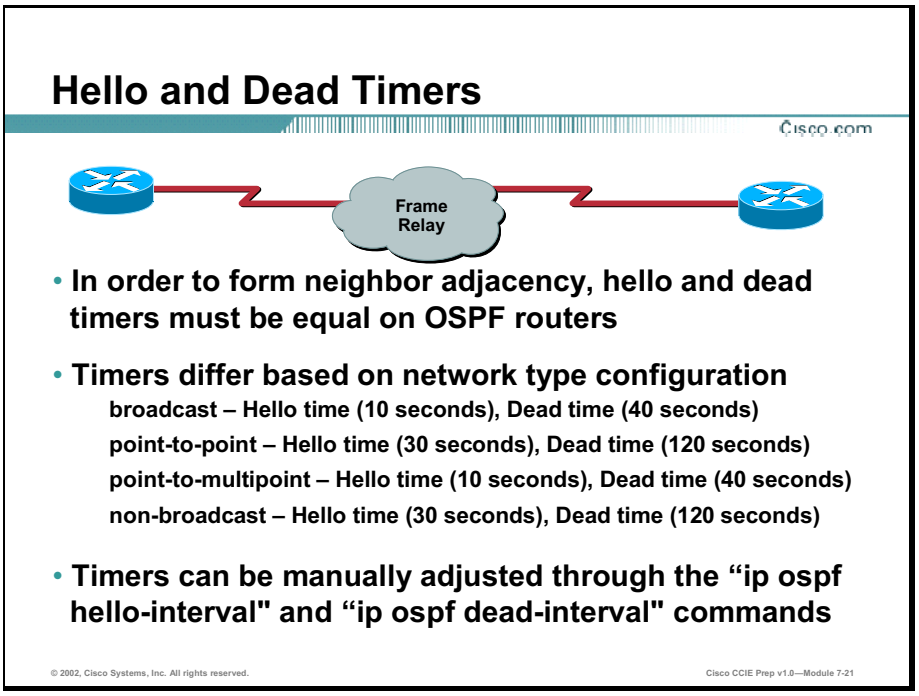

OSPF requires these intervals to be exactly the same between two routers in order for the routers to form a neighbor adjacency. If either of these intervals is different, the routers will not become neighbors on a particular segment. Changing the OSPF network type with the **ip ospf network** command affects these intervals. Here is a list of the default intervals for the different network types.

- **Broadcast:** Hello time 10 seconds, Dead time 40 seconds
- **non-broadcast:** Hello time 30 seconds, Dead time 120 seconds
- **point-to-point:** Hello time 10 seconds, Dead time 40 seconds
- **point-to-multipoint:** Hello time 30 seconds, Dead time 120 seconds

In some situations you may need to manually set these timers on one router in order for that router to form a neighbor adjacency with another router. The interface configuration commands used to set these timers are shown below.

Table 6-5: OSPF Timer Commands

Command	Description
`ip ospf hello interval` *`seconds`*	Manually sets the hello timer on an interface. By default, the dead timer is set to four times the hello timer
`ip ospf dead interval` *`seconds`*	Manually sets the dead timer on an interface

OSPF Operation in an NBMA Topology

This section discusses OSPF in a Non-Broadcast Multi-Access (NBMA) environment.

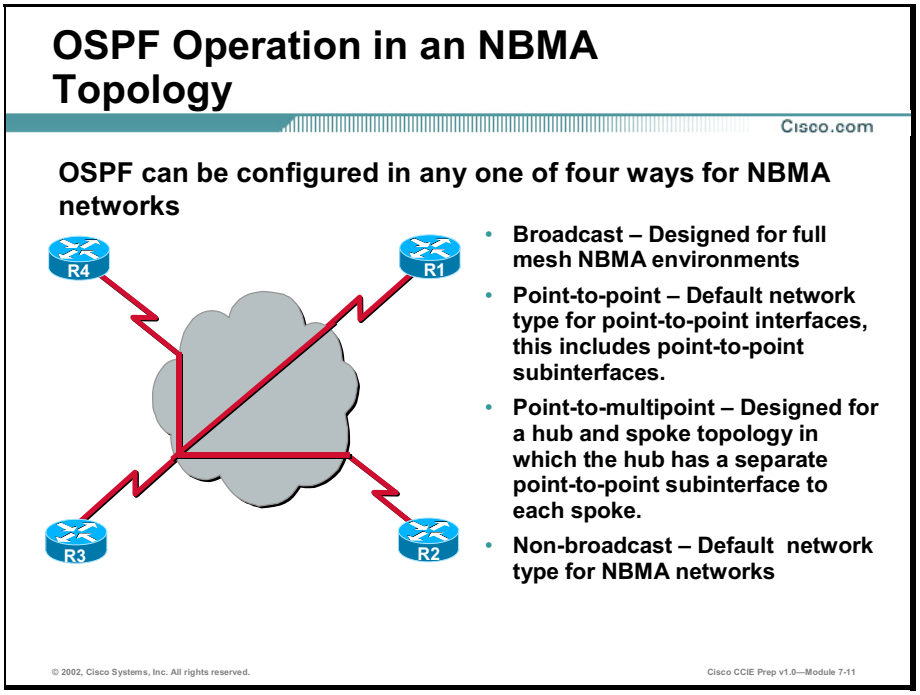

Special care should be taken when configuring OSPF over NBMA networks, such as Frame Relay, X.25, and Asynchronous Transfer Mode (ATM). By default, OSPF treats these media types like any other broadcast media (Ethernet, Token Ring, etc.). This causes problems because most NBMA networks are usually configured in a partial mesh or hub and spoke topology, where all routers are not directly connected to one another.

To simplify the configuration of OSPF over Frame Relay on Cisco routers, use the **ip ospf network** command. This command allows you to control what type of network OSPF thinks it is dealing with. Four different network types can be defined using this command.

- **broadcast** – Designed for full mesh NBMA environments.
- **point-to-point** – Default network type for point-to-point interfaces, this includes point-to-point subinterfaces.
- **point-to-multipoint** – Designed for a hub and spoke topology in which the hub has a separate point-to-point subinterface to each spoke. Can also be used when the hub uses a physical interface or point-to-multipoint subinterface to communicate with spoke routers.
- **non-broadcast** – Default network type for NBMA networks.

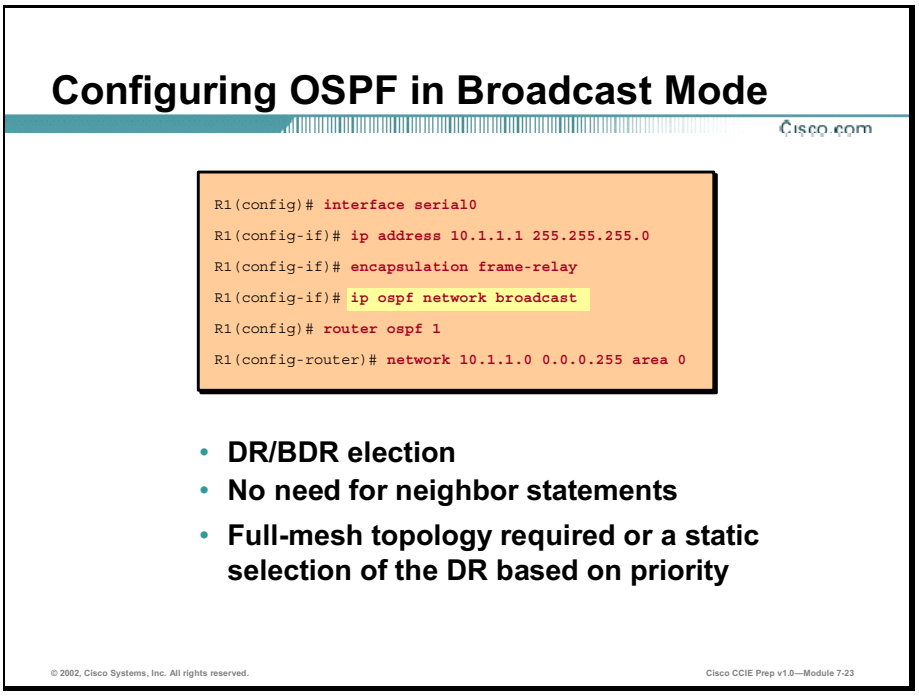

If your NBMA network is configured in a full mesh, there is no need to statically define neighbors, as all routers can reach each other directly. This also eliminates the need to carefully control the DR/BDR election process. For the broadcast network type to work successfully, the **broadcast** parameter must be specified on all Frame Relay map statements.

Table 6-6: < ip ospf network broadcast > Command

Command	Description
`ip ospf network broadcast`	Used for full mesh NBMA networks. Does not require neighbors to be statically defined. Requires the election of a DR/BDR. Performed from router configuration mode.

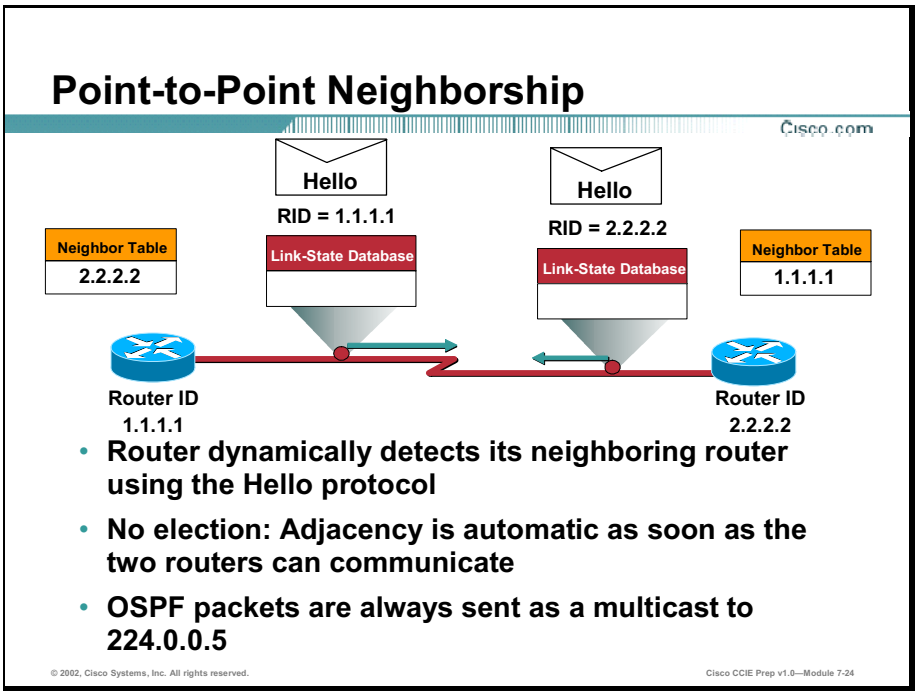

A point-to-point network connects a single pair of routers. A T1 serial line is a common example of a point-to-point network.

On point-to-point networks, OSPF routers dynamically detect their neighboring routers by sending hello packets to the All OSPF Routers multicast address of 224.0.0.0.5. Since there are only two routers on a point-to-point network there is no need for a DR/BDR election. On a point-to-point network, OSPF routers become neighbors as soon as they see themselves in the other router's hello packet.

Usually, the source address of an OSPF packet is set to the Internet Protocol (IP) address of the outgoing interface on the router. It is possible however to use IP unnumbered interfaces with OSPF. On unnumbered interfaces, the source address will be set to the IP address of the interface that the unnumbered interface is borrowing its IP address from.

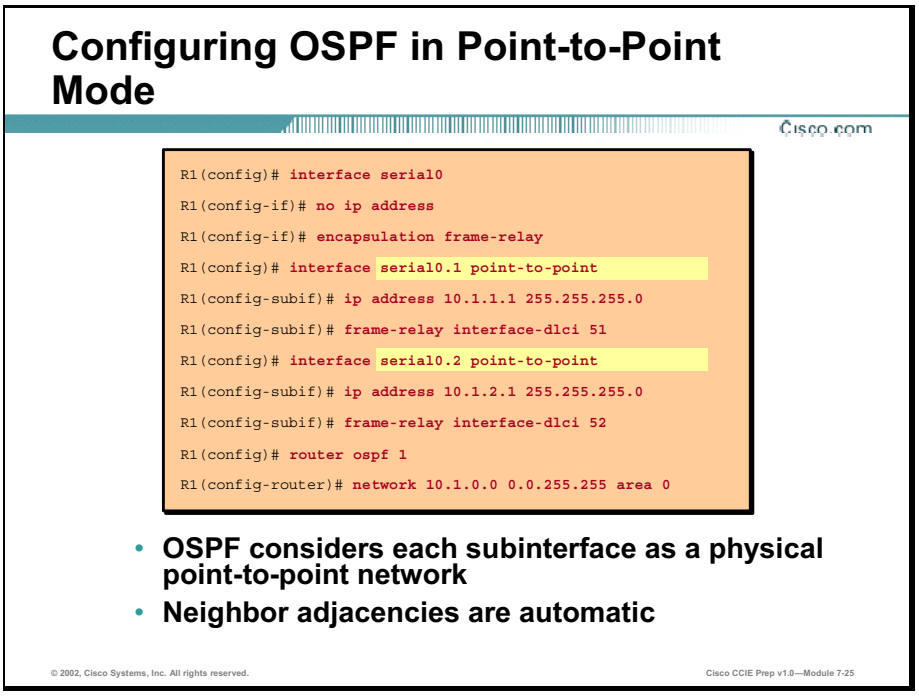

Point-to-point subinterfaces on the hub router treat each connection to a spoke router as a separate IP subnet. Point-to-point subinterfaces were originally created in order to handle issues caused by split horizon when running distance vector routing protocols over NBMA networks.

A point-to-point subinterface has the same properties of a physical point-to-point interface. As far as OSPF is concerned, an adjacency is always formed over a point-to-point subinterface, with no DR/BDR election.

Point-to-point mode is the default OSPF network type for point-to-point subinterfaces. Therefore, no further configuration is required.

Table 6-7: < ip ospf network point-to-point > Command

Command	Description
`ip ospf network point-to-point`	Sets the OSPF network type to point-to-point. Performed from router configuration mode.

Note By default, OSPF considers loopbacks as host routes and advertises them with a /32 subnet mask. You can add the **ip ospf network point-to-point** command to a loopback interface and OSPF will then advertise the loopback interface's actual subnet mask. This is useful when classless-to-classful route redistribution is being performed in the network. Route redistribution is covered in Module 9: Advanced Routing Techniques.

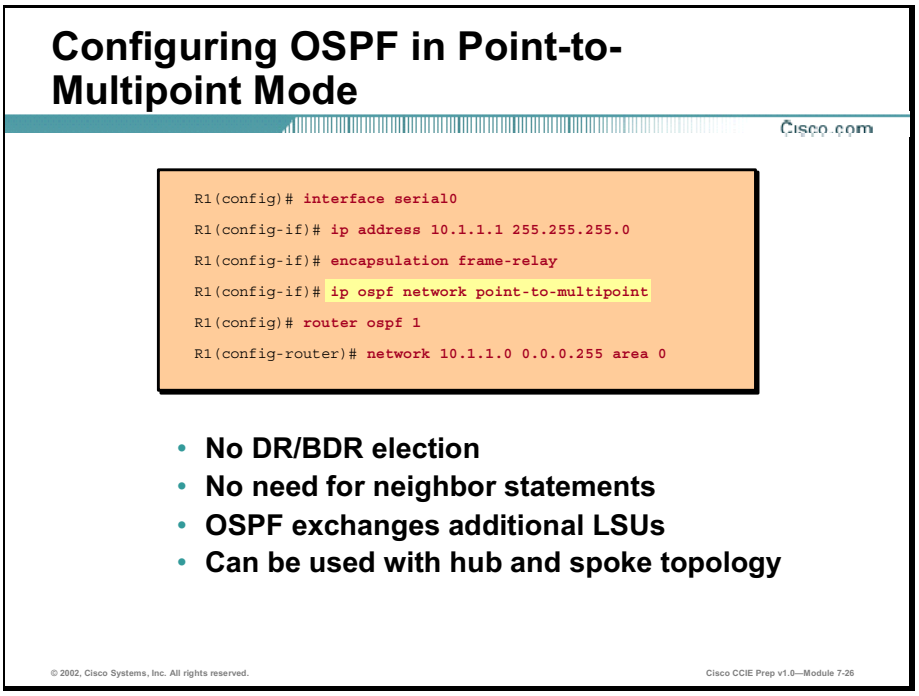

To avoid the headaches with non-broadcast and broadcast modes, the point-to-multipoint network type is available. An OSPF point-to-multipoint network is seen as one or more numbered point-to-point interfaces. As with the non-broadcast and broadcast modes, the NBMA cloud is seen as one IP subnet. The main advantage to using the point-to-multipoint network type is that it does not require the use of a DR/BDR. OSPF point-to-multipoint networks avoid this by exchanging additional link-state updates that contain a number of information elements that describe connectivity to the neighboring routers.

Point-to-multipoint OSPF networks can be configured as either broadcast or non-broadcast. Broadcast networks are configured with the **ip ospf network point-to-multipoint** command. This mode is RFC-compliant and functions exactly like the broadcast network type without the need for a DR/BDR. Non-broadcast point-to-multipoint networks are configured with the **ip ospf network point-to-multipoint non-broadcast** command. The non-broadcast network is a Cisco extension and functions exactly like a non-broadcast network type, where neighbors are statically defined. However, there is still no need for a DR/BDR.

Table 6-8: < ip opsf network point-to-multipoint [non-broadcast]> Command

Command	Description
`ip opsf network point-to-multipoint [non-broadcast]`	Functions like a **broadcast** network, without the need for a DR/BDR. Performed from router configuration mode. **non-broadcast** – Functions like a non-broadcast network, without the need for a DR/BDR. Requires OSPF neighbors to be statically defined using the **neighbor** command.

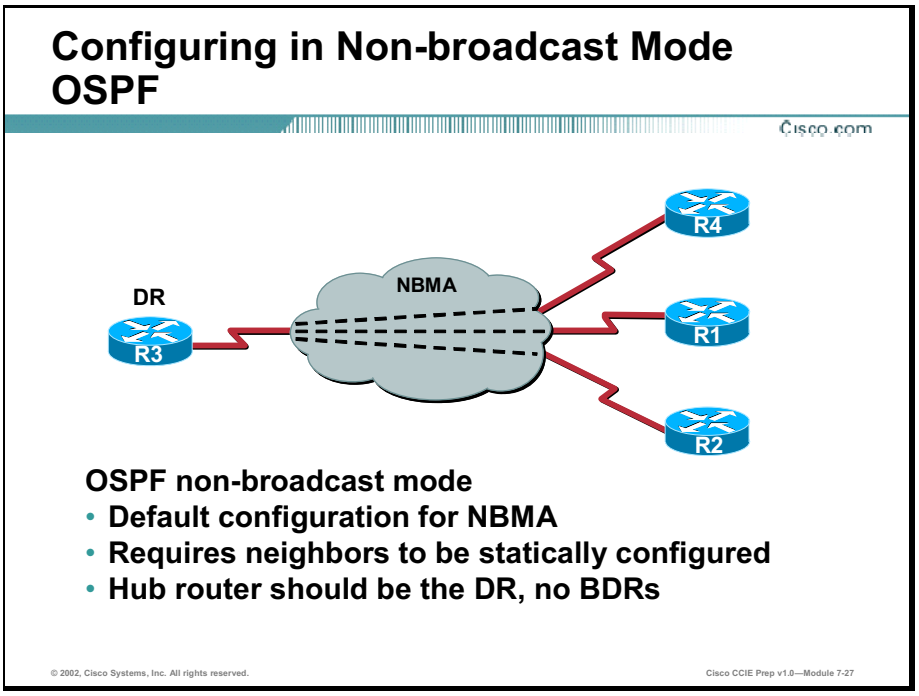

When the network type is set to non-broadcast, which is the default on NBMA networks, OSPF operates as if it were running in a broadcast multi-access environment, such as Ethernet. Therefore, a DR and BDR are elected for the NBMA network, and the DR originates the Link-State Advertisements (LSAs) for the network. If you are operating in a full mesh environment and the **broadcast** keyword is specified in your Frame Relay map statements, no other configuration is needed. However, in a hub and spoke topology, OSPF neighbors must be statically configured using the **neighbor** command.

The **ip ospf priority** command must be set to 0 on all spoke routers as well to ensure that the hub router becomes the DR and there is no BDR elected. The hub router is required to become the DR, since it is the only router that has full connectivity to all other routers in the network. This restriction also requires that no BDR be elected.

For example, if one of the spokes is acting as a BDR and the hub router, which is performing the DR function, goes down, one of the spoke routers will take over as the DR. The spoke router will not have connectivity to all other routers in the network and communication will fail. Also, when the hub router comes back online it will not resume the role of the DR, since the spoke router has already attained that role and is still functioning.

Table 6-9: < ip ospf network non-broadcast > Commands

Command	Description
`ip ospf network non-broadcast`	The default network type for NBMA networks. Requires statically defined neighbors and the use of a DR/BDR.

Command	Description
neighbor ip address *priority <0-255> poll-interval <sec> cost <1-65535>*	Router configuration command that statically defines the router's OSPF neighbors. **ip address** – IP address of the neighbor. **priority** – (Optional) 8-bit number that indicates the priority value (used during the DR/BDR election) of the non-broadcast neighbor. The default is 0. Neighbors with no specific priority configured will assume the priority assigned to their interface that connects to the NBMA network. **poll-interval** - (Optional) If a neighboring router has become inactive (hello packets have not been seen and the dead interval has elapsed), it may still be necessary to send hello packets to the dead neighbor. These hello packets will be sent at a reduced rate called the poll interval. RFC1247 recommends that this value be much larger than the hello interval. The default is 120 seconds (2 minutes). **cost** – (Optional) Assigns a cost (1-65535) to the neighbor. Neighbors with no specific cost configured will assume the cost of their interface, based on the bandwidth or the **ip ospf cost** command.

OSPF over NBMA Topology Summary

Cisco.com

Mode	Preferred Topology	Subnet Address	Adjacency
Non-broadcast	Fully meshed	Same	Manual configuration DR/BDR elected
Broadcast	Fully meshed	Same	Automatic DR/BDR elected
Point-to-multipoint	Partial mesh (hub and spoke)	Same	Automatic No DR/BDR
Point-to-multipoint nonbroadcast	Partial mesh (hub and spoke)	Same	Manual configuration No DR/BDR
Point-to-point	Partial mesh (hub and spoke), using subinterfaces	Different for each subint.	Automatic No DR/BDR

This table provides a concise comparison of the different modes of operation for OSPF over NBMA topologies.

Summary

This section summarizes the key points discussed in this lesson.

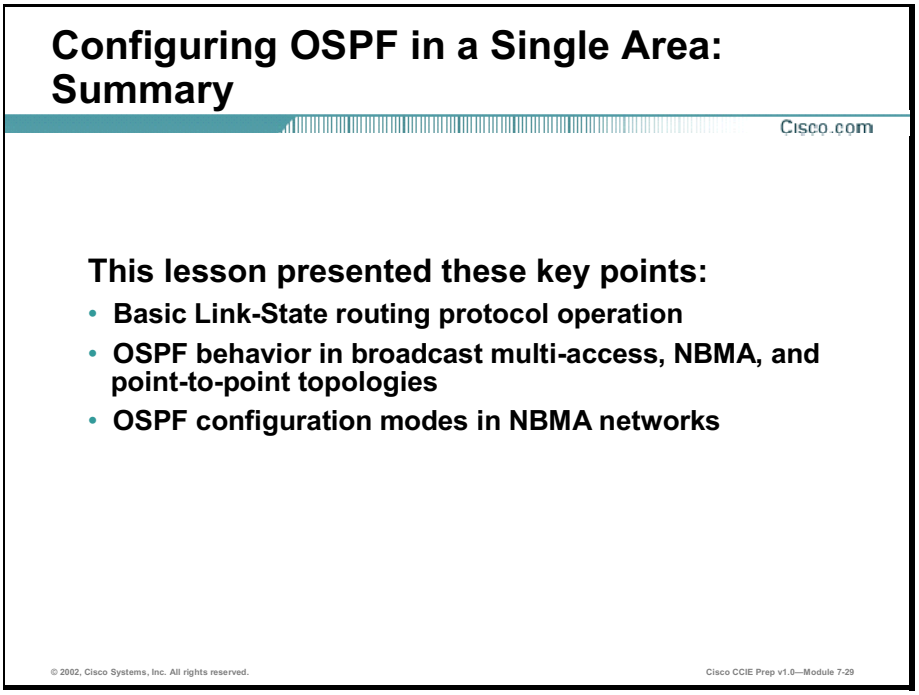

Next Steps

After completing this lesson, go to:

- Multi-area OSPF Environments

References

For additional information, refer to these resources:

- http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121cgcr/ip_c/ipcprt2/1cdospf.htm

Lesson Assessment (Quiz)

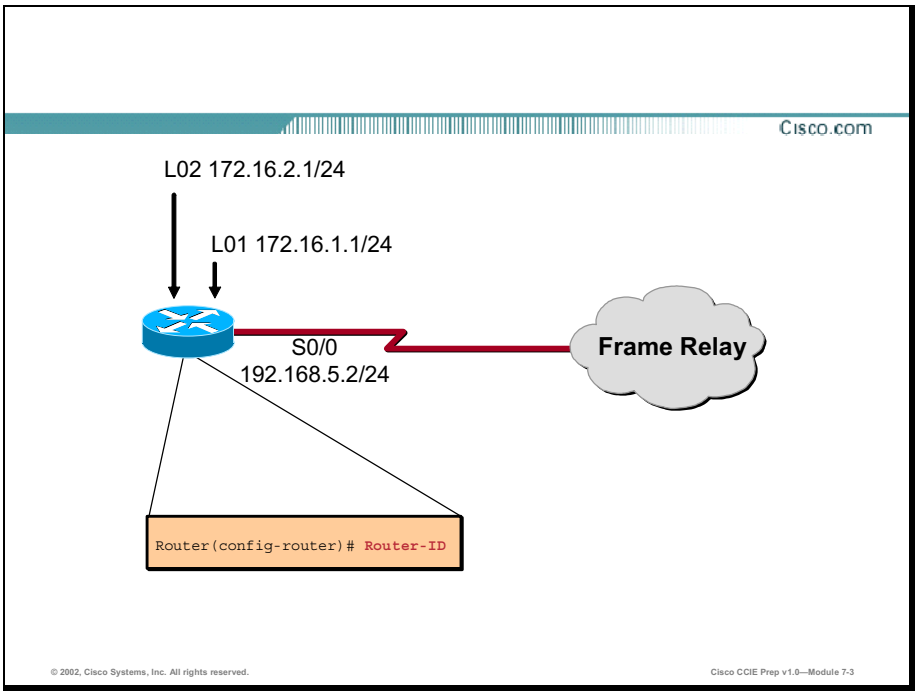

Q1) Based on the configuration above, what will the router ID of this router be?

Q2) Which of the following OSPF priority values is used to prevent a router from participating in the DR/BDR election?

A) 0

B) 1

C) 255

D) There is no way to prevent a router from participating in the DR/BDR election

Q3) What command is used to prevent Fast Ethernet and Gigabit Ethernet from both having an OSPF cost of 1?

Q4) Which OSPF network type requires statically defined neighbors and strict control of the DR/BDR election in a hub and spoke NBMA topology?

A) broadcast

B) non-broadcast

C) point-to-point

D) point-to-multipoint

Q5) Which OSPF network types do not require a DR/BDR election?

A) broadcast

B) non-broadcast

C) point-to-point

D) point-to-multipoint

Multi-Area OSPF Environments

Overview

Single area Open Shortest Path First (OSPF) deployments can become hard to manage when hundreds of routers are involved. One solution is to break up the OSPF domain into multiple areas, allowing for route summarization and a hierarchical routing structure. This lesson examines key criteria to be used when deciding to divide OSPF networks into multiple areas and the necessary OSPF configuration to accomplish this.

Importance

OSPF is the core Interior Gateway Protocol (IGP) used in the Cisco Certified Internetworking Expert (CCIE) lab.

Objectives

Upon completing this lesson, you will be able to:

- Configure OSPF in a multi-area environment
- Describe the use of stub areas and the differences between stub, totally stubby, and not-so-stubby areas
- Configure inter-area and external route summarization

Learner Skills and Knowledge

To fully benefit from this lesson, you must have these prerequisite skills and knowledge:

- Passed the CCIE written qualification exam
- Completed the Building Scalable Cisco Internetworks (BSCI) course or have the equivalent knowledge

Outline

This lesson includes these sections:

- Overview
- Configuring OSPF in a Multi-area Environment
- Route Summarization
- Summary

Configuring OSPF in a Multi-Area Environment

This section describes OSPF configuration in a multi-area environment.

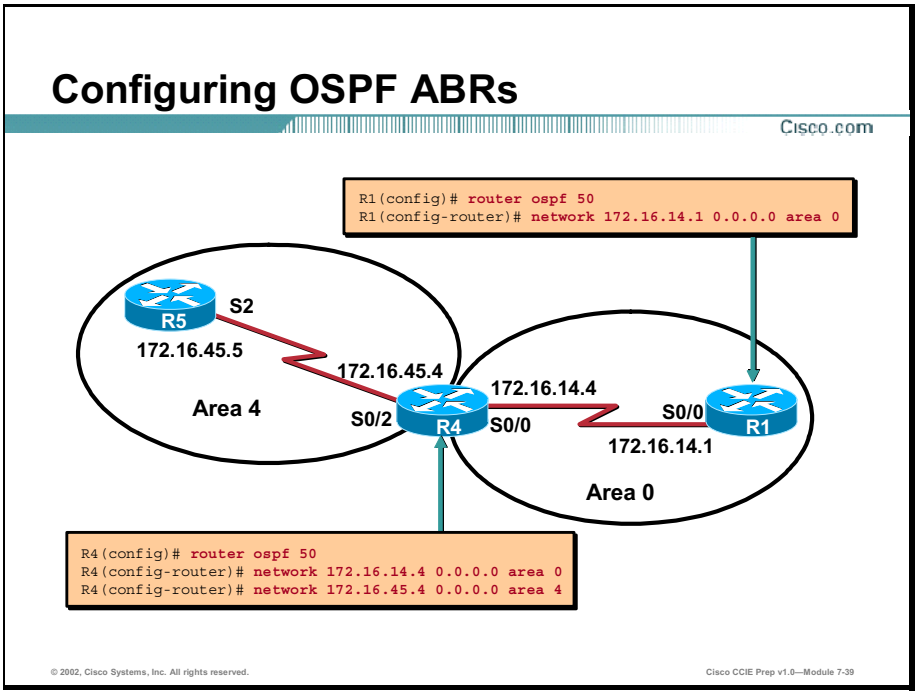

There are no special commands to make a router an Area Border Router (ABR) or Autonomous System Boundary Router (ASBR). The router assumes this role by virtue of the areas or autonomous systems to which it is connected. If a router has interfaces in different OSPF areas, it is an ABR. If the router has an interface in OSPF and an interface in another routing protocol and is performing redistribution of that protocol into OSPF, it is an ASBR.

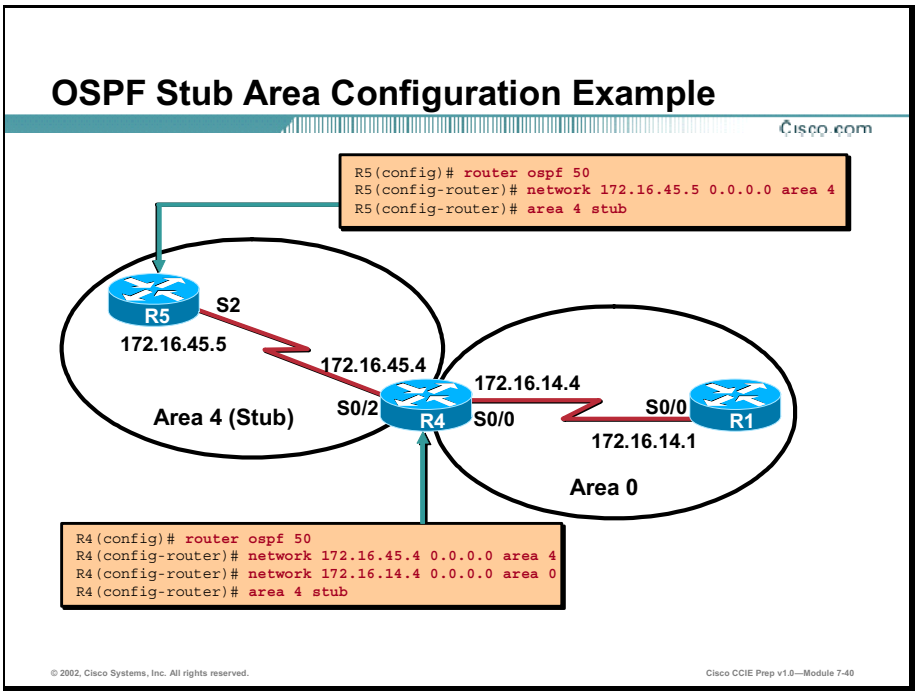

In the example above, Area 4 is configured as a stub area. No external routes (Type 5 Link-State Advertisements (LSAs)) from the external autonomous systems (other routing domains) will be allowed into the stub area.

The **area 4 stub** command on each router in Area 4 defines the stub area. Each router in the stub area must be configured with the **area 4 stub** command or neighbor adjacencies will not be formed.

Table 6-10: < area <*area-id*> stub > Command

Command	Description
area <*area-id*> **stub**	Designates this router as part of a stub area. Performed in router configuration mode.

The only routes that will appear in R4's routing table are intra-area routes (designated with an O in the routing table), inter-area routes, and the default route (these routes will be designated with an O-IA in the routing table).

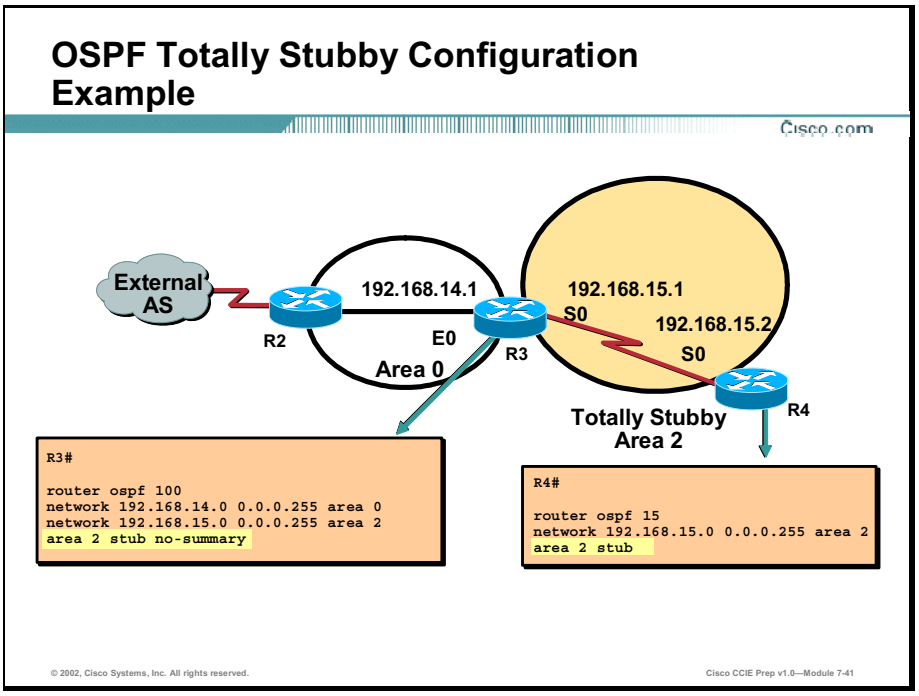

In this example, the keyword **no-summary** has been added to the **area 2 stub** command on R3. This keyword causes summary routes (inter-area Type 3/4 LSAs) to also be blocked from the stub area. Each router in the totally stubby area will pick the closest ABR as a gateway to get to any destinations outside of the area. The only routes that will appear in R4's routing table now are intra-area routes (designated with an O in the routing table) and the default route. No inter-area routes (designated with O-IA in the routing table) will be shown, except the default route.

Table 6-11: < area <*area-id*> stub no-summary > Command

Command	Description
area <*area-id*> **stub no-summary**	Prevents an ABR from sending summary link advertisements into the stub area. Use this keyword to create a totally stubby area. Performed in router configuration mode.

Note The **no-summary** keyword is only required on the ABRs connected to the totally stubby area. The internal routers within the stub area only need to be configured with the **area stub** command.

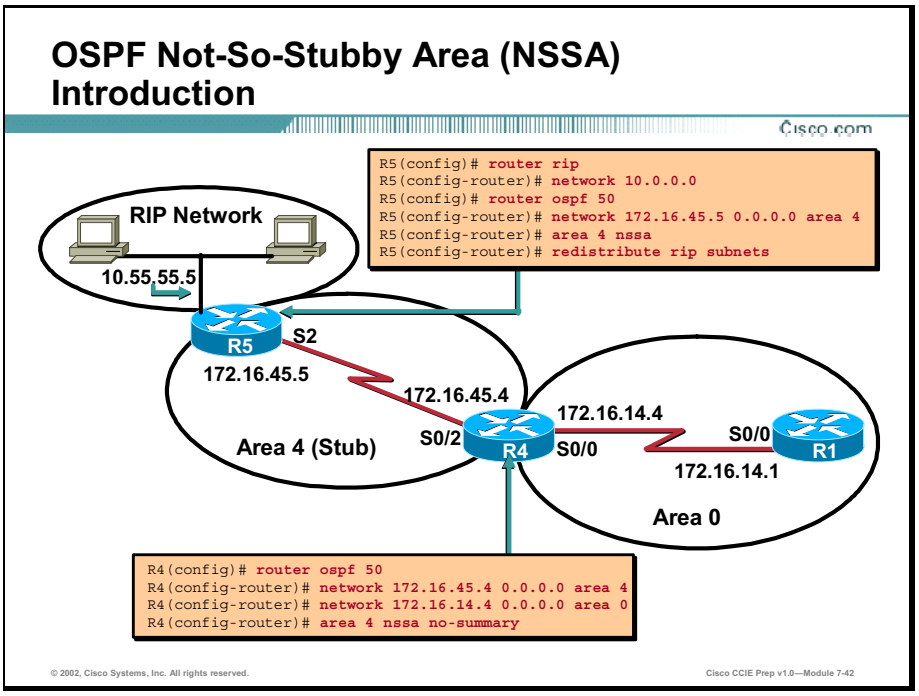

In the network diagram above, Area 4 is defined as a stub area. R5 also connects to a Routing Information Protocol (RIP) network and therefore qualifies as an ASBR. However, the RIP routes cannot be propagated into the OSPF domain, because Type 5 LSAs are not allowed into a stub area. In order to redistribute RIP information into OSPF here, we must define area 4 as a special type of stub area called a not-so-stubby area (NSSA).

Since Type 5 LSAs are not allowed into a stub area, the NSSA ASBR generates Type 7 LSAs instead, which are flooded throughout the NSSA. The Type 7 LSAs get translated into Type 5 LSAs by the NSSA ABR and are propagated into the OSPF backbone. Redistributed Enhanced Interior Gateway Routing Protocol (EIGRP) and Interior Gateway Routing Protocol (IGRP) routes from other areas in the network are still not allowed into area 4, because a NSSA is just an extension to the stub area, and all the characteristics of a stub area still exist, including disallowing Type 5 LSAs.

Table 6-12: < area <*area-id*> nssa no-summary > Command

Command	Description
area <*area-id*> **nssa no-summary**	Configures a stub area that contains an ASBR to allow external routes into the stub area as Type 7 LSAs. The **no-summary** keyword is used to create a totally stubby NSSA. Performed in router configuration mode.

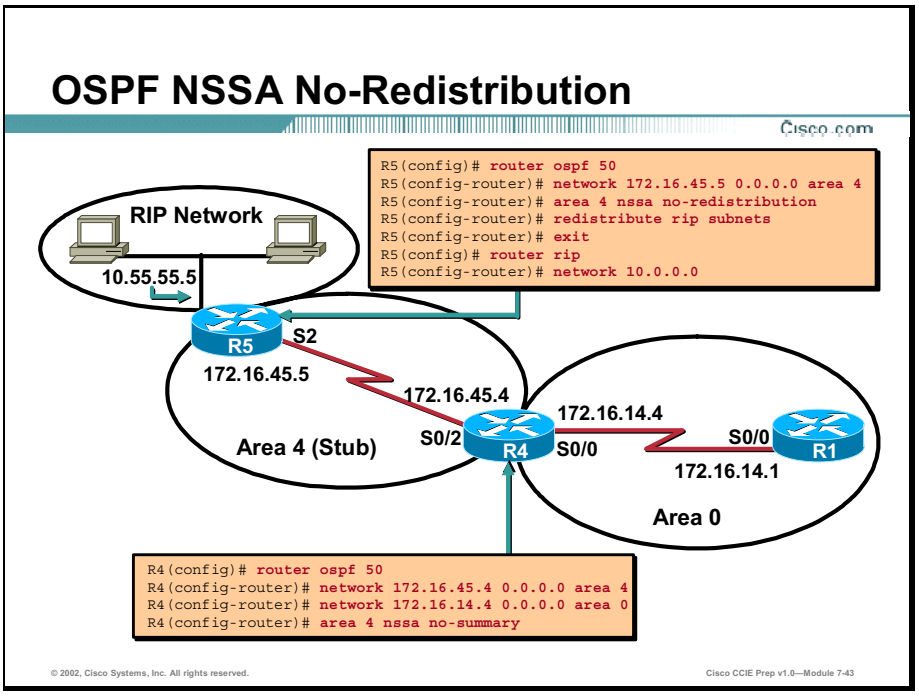

NSSA no-redistribution can be used to stop redistributed Type 7 routes from being flooded in the NSSA area.

Further, no-summary can be added to stop Type 3/4 LSAs from being distributed except for the default route Type 3 summary.

Table 6-13: < area <*area-id*> nssa no-redistribution > Command

Command	Description
`no-redistribution`	(Optional) Used when the router is an NSSA Area Border Router (ABR) and you want the **redistribute** command to import routes only into the normal areas, but not into the NSSA area. Performed in router configuration mode.

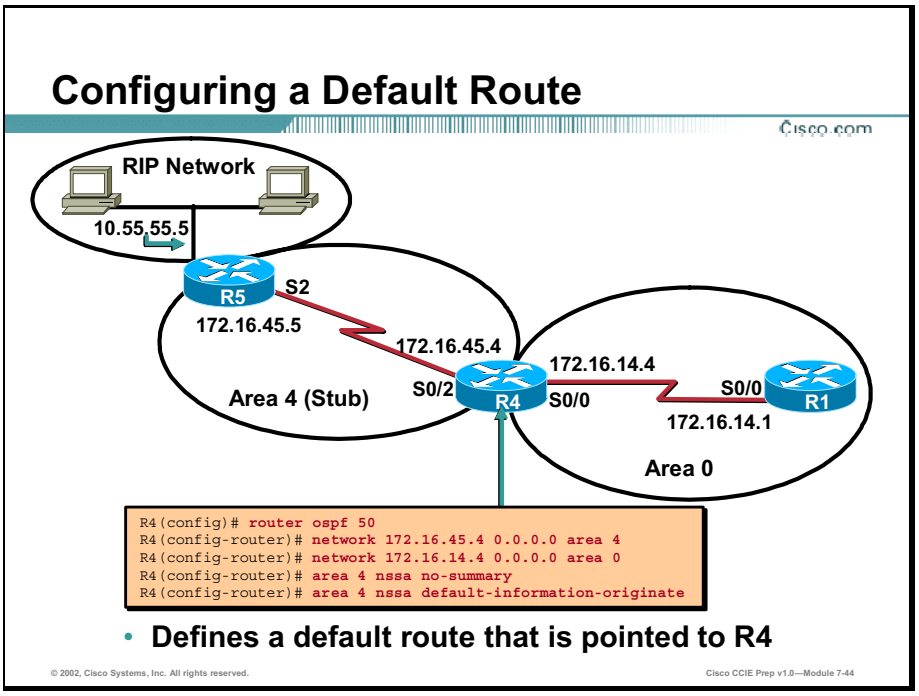

NSSA default-information originate command will allow Type 3/4 LSAs into the area and define a default route that is pointed to the ASBR, R4.

This command can be preceded by the no-redistribution to stop Type 7 LSAs from being propagated.

Table 6-14: < area <*area-id*> nssa default-information-originate > Command

Command	Description
`default-information-originate`	(Optional) Used to generate a Type 7 default into the NSSA area. This keyword only takes effect on NSSA ABRs or NSSA Autonomous System Boundary Routers (ASBRs). Performed in router configuration mode.

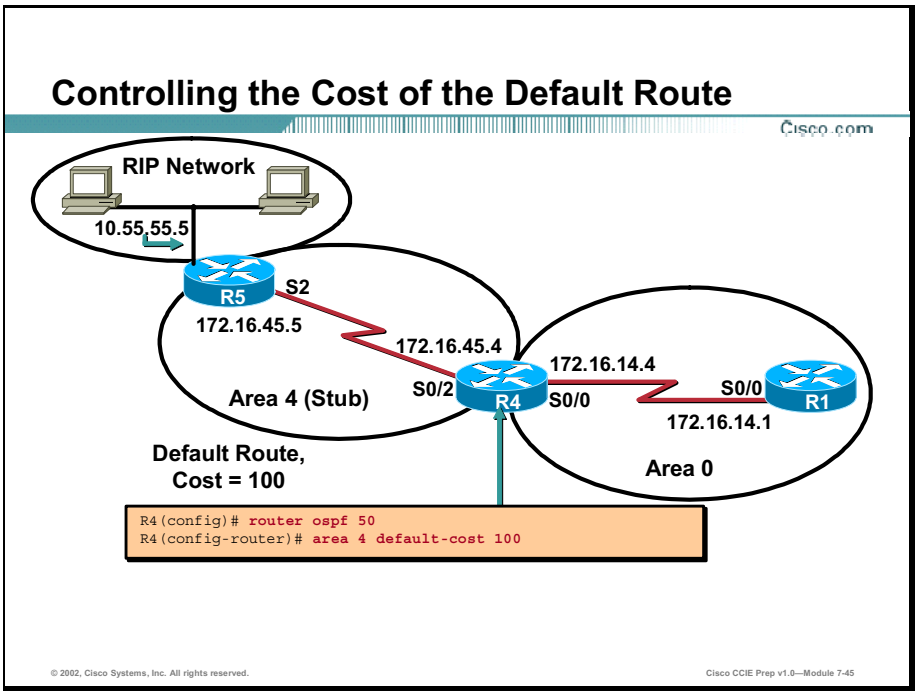

By default, the default route injected into the stub area by the ABR has a cost of 1. This cost will increment as the default route is propagated throughout the stub area. There may be instances where you will want the default route to start with a higher default cost than 1.

Table 6-15: < area <*area-id*> default-cost <0-16777215>> Command

Command	Description
area <*area-id*> **default-cost** <0-16777215>	Defines the starting cost of the default route that is injected into the stub, totally stubby, or not-so-stubby area. Performed in router configuration mode. (Optional command for ABRs only)

Route Summarization

This section covers route summarization.

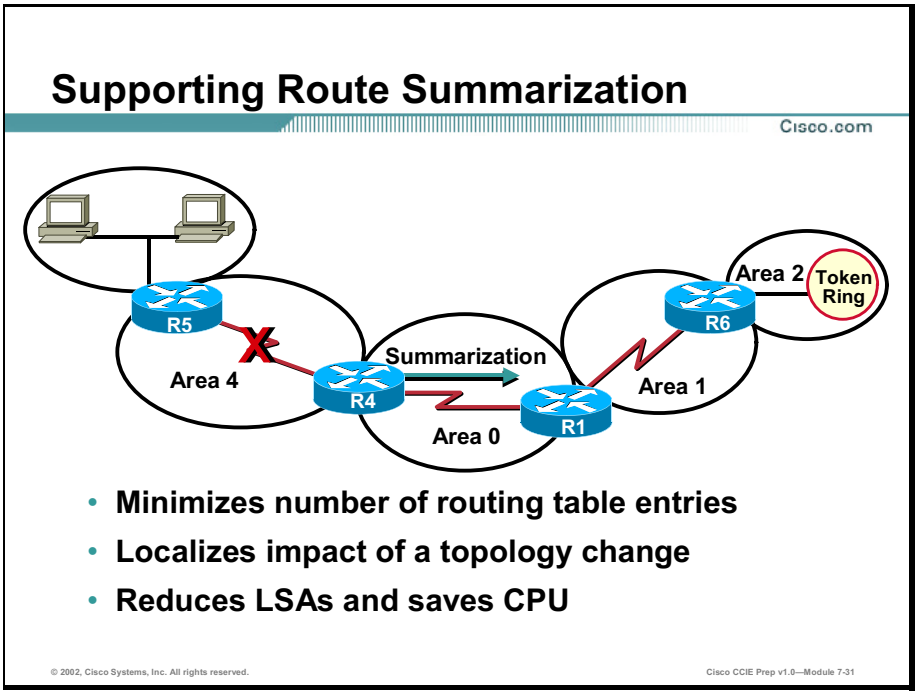

With route summarization, only summarized routes will be propagated into the backbone (Area 0). This is very important because it prevents every router in the OSPF domain from having to rerun the Shortest Path First (SPF) algorithm when a route changes within an area. This increases the network's stability and reduces unnecessary traffic.

There are two types of summarization:

- **Inter-area route summarization:** Inter-area route summarization is done on ABRs and applies to routes from within each area. It does not apply to external routes injected into OSPF via redistribution. To perform effective route summarization, network numbers within the areas should be assigned in a contiguous fashion, so that these addresses can be summarized into a minimal number of summary addresses.

- **External route summarization:** External route summarization is specific to external routes that are injected into OSPF via redistribution. Again, it is important to ensure that the external address ranges being summarized are contiguous. Summarizing overlapping ranges from two different routers could cause packets to be sent to the wrong destination. ASBRs are usually the only routers that perform external route summarization.

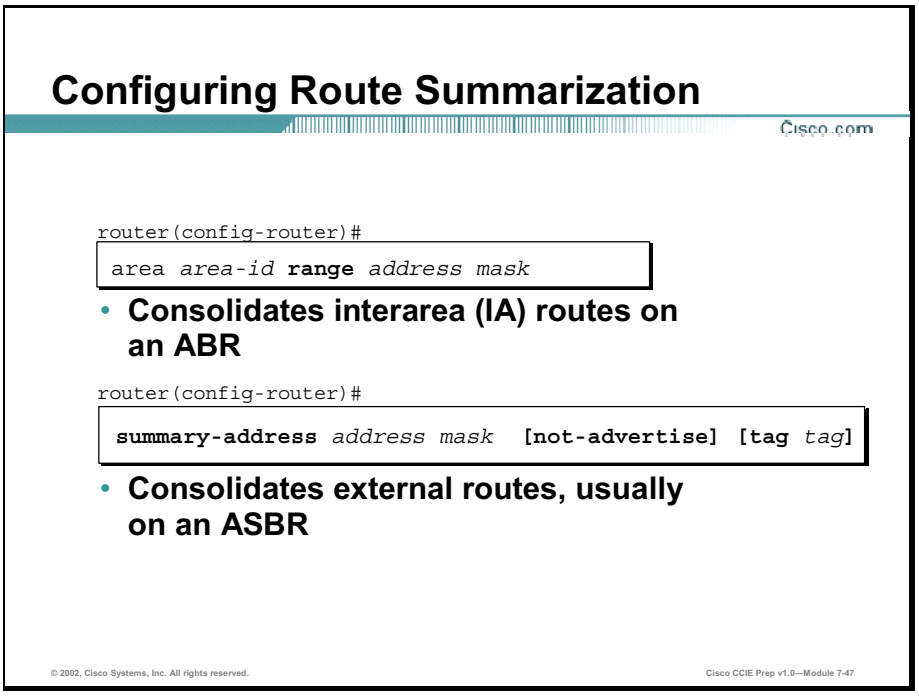

OSPF does not perform auto summarization. To configure manual inter-area route summarization, use the **area range** command. This command instructs the ABR to summarize routes for a specific area before injecting them into the backbone.

Table 6-16: < area <*area-id*> range *summary address summary mask* > Command

Command	Description
area <area> range *summary address summary mask*	*area* - Identifies the area from which routes are to be summarized *summary address* - Summary address for a range of addresses *summary mask* - Subnet mask used to summarize the more specific routes into one advertisement

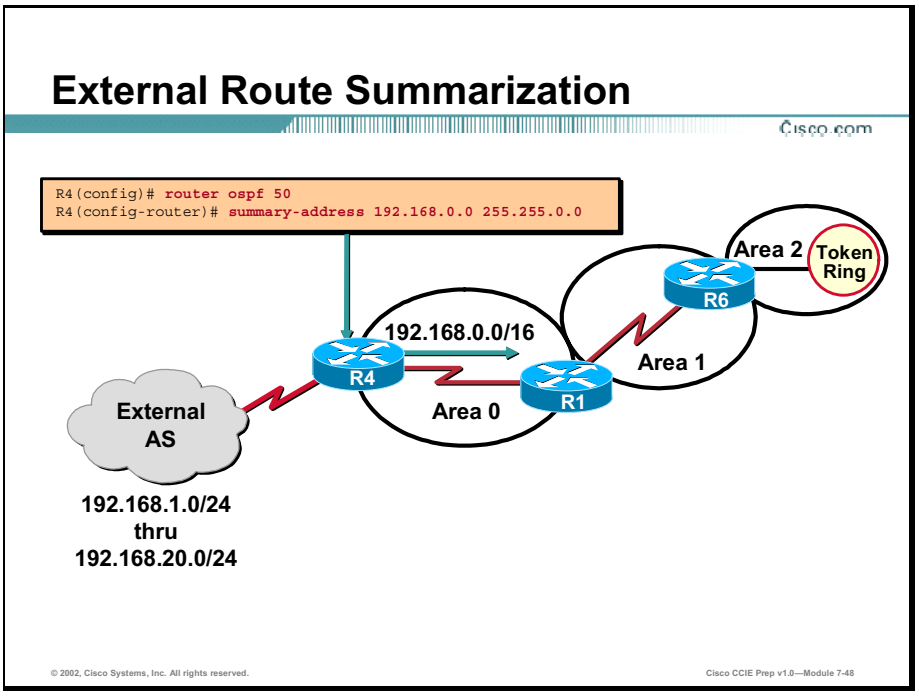

To configure manual route summarization on an ASBR to summarize external routes, use the **summary-address** command. This command instructs the ASBR to summarize external routes before injecting them into the OSPF domain.

Table 6-17: < summary-address > Command

Command	Description
summary-address *summary address summary mask*	Summary address designated for a range of external routes. Performed in router configuration mode. *summary address* - summary address for a range of external addresses. *summary mask* - subnet mask used to summarize the more specific routes into one advertisement.

Note The **summary-address** command can also be used on ASBRs to summarize routes within Area 0. This technique is extremely useful when performing classless-to-classful route redistribution.

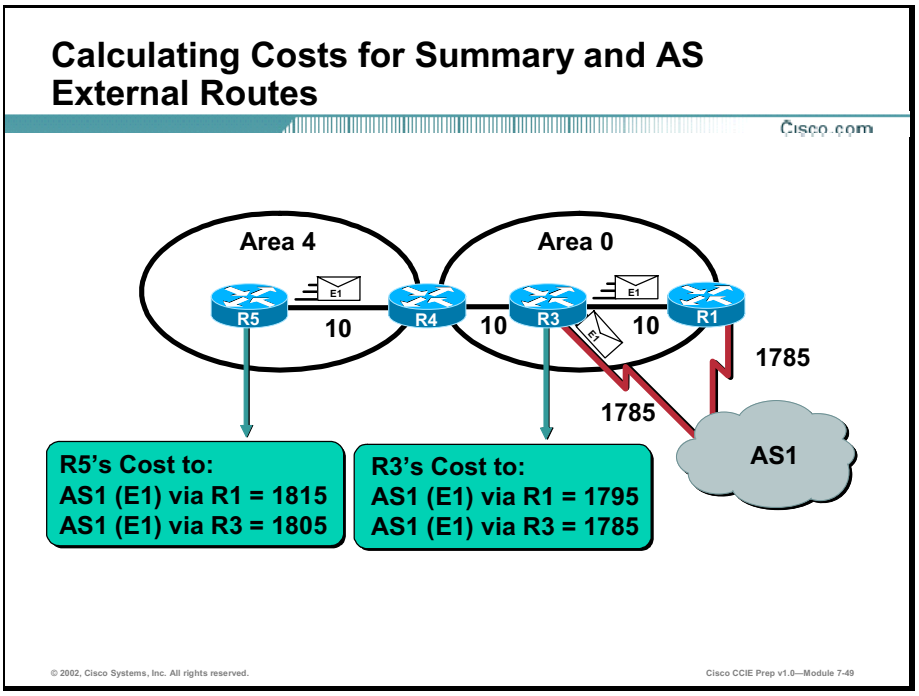

This page discusses the process of calculating the cost for summary and external routes:

- Calculating the cost of summary routes
  - The cost of a summary route is the smallest cost of a given inter-area route that appears in the summary plus the cost of the ABR link to the backbone. For example, if the cost of the ABR link to the backbone is 50, and the smallest cost of a route being summarized is 49, the total cost associated with the summary route would be 99. This cost is calculated automatically for each summary route.
- Calculating the cost of external routes
  - The cost of an external route differs depending on the metric-type set for external route on the ASBR. The metric-type is set during route distribution. The default metric-type is E2.
    - **Type 1 (E1):** If an external route has a metric-type of E1, then the metric is calculated by adding the external cost (configured during redistribution) to the internal cost of each link that the route advertisement crosses in the OSPF domain. Use this metric-type when multiple ASBRs are advertising the same external routes an AS.
    - **Type 2 (E2)** (default) If an external route has a metric-type of E2, it will always have the external cost (configured during redistribution) assigned. The metric does not increment as the route advertisement passes throughout the OSPF domain. Use this metric-type if only one ASBR is advertising external routes to the AS. Type 2 routes are preferred over type 1 routes unless two equal-cost routes exist to the external destination.

Summary

This section summarizes the key points discussed in this lesson.

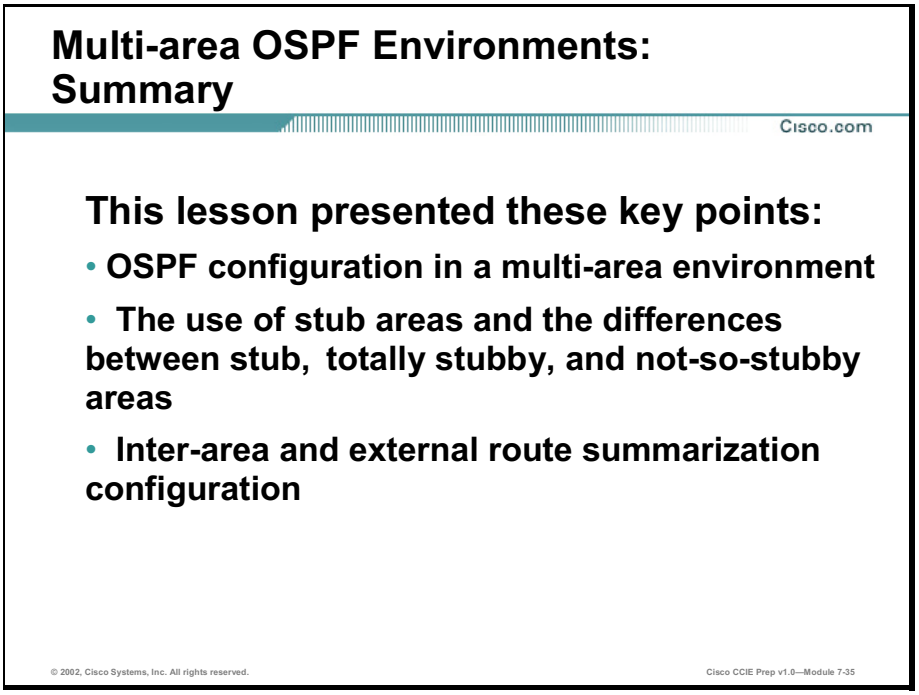

Next Steps

After completing this lesson, go to:

- Advanced OSPF Features

References

For additional information, refer to these resources:

- OSPF Design Guide - http://www.cisco.com/warp/public/104/1.html

Lesson Assessment (Quiz)

Q1) List the different types of stub areas that Cisco routers support.

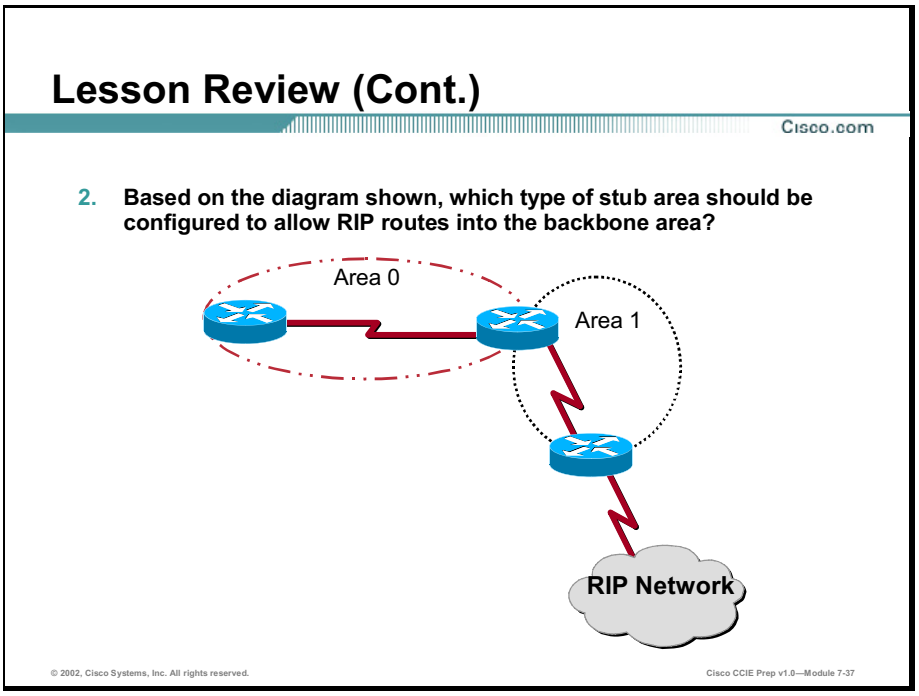

Q2) Based on the diagram above, which type of stub area should be configured to allow RIP routes into the backbone area?

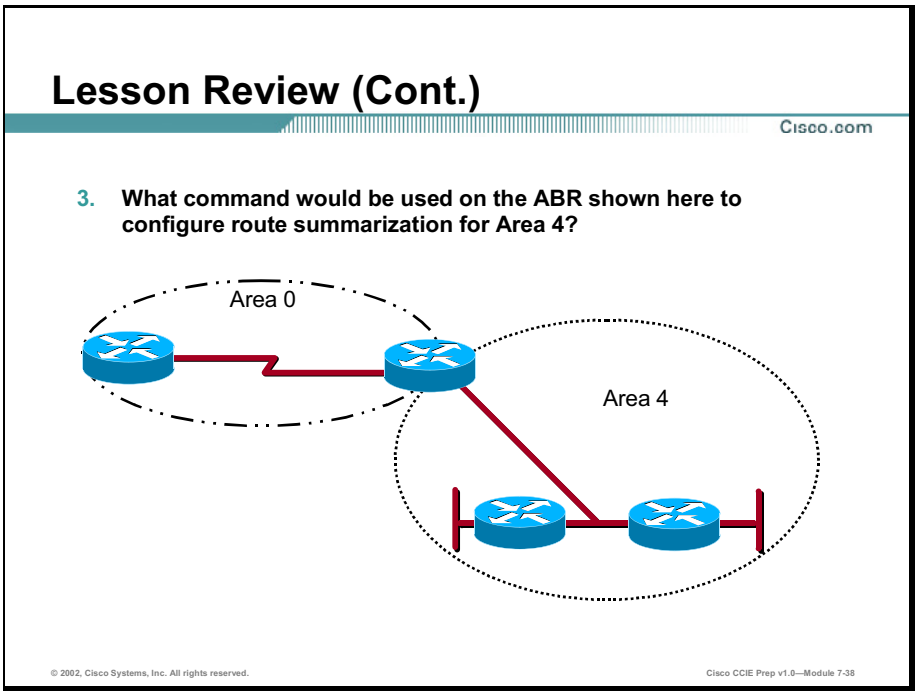

Q3) What command would be used on the ABR shown above to configure route summarization for Area 4?

Q4) What command is used to configure external route summarization on an ASBR?

Q5) What type of external route increments its cost as it is propagated throughout the OSPF domain?

Advanced OSPF Features

Overview

In some situations such as company mergers or buyouts, the standard Open Shortest Path First (OSPF) configurations do not allow for an easy migration of routing protocols. This lesson examines many of the advanced OSPF features that can be used to expand, secure, and optimize your OSPF network.

Importance

OSPF is the core Interior Gateway Protocol (IGP) used in the Cisco Certified Internetworking Expert (CCIE) lab.

Objectives

Upon completing this lesson, you will be able to:

- Configure virtual links in an OSPF multi-area environment
- Configure OSPF neighbor authentication
- Configure OSPF demand circuits to prevent OSPF hellos from bringing up Integrated Services Digital Network (ISDN) Dial-on-Demand Routing (DDR) links

Learner Skills and Knowledge

To fully benefit from this lesson, you must have these prerequisite skills and knowledge:

- Passed the CCIE Written qualification exam
- Completed the Building Scalable Cisco Internetworks (BSCI) course or have the equivalent knowledge

Outline

This lesson includes these sections:

- Overview
- Virtual Links Overview
- OSPF Authentication
- OSPF Demand Circuits
- Summary
- Lesson Assessment (Quiz)

Virtual Links Overview

This section describes the function of virtual links and discusses their relevance in a multi-area OSPF environment.

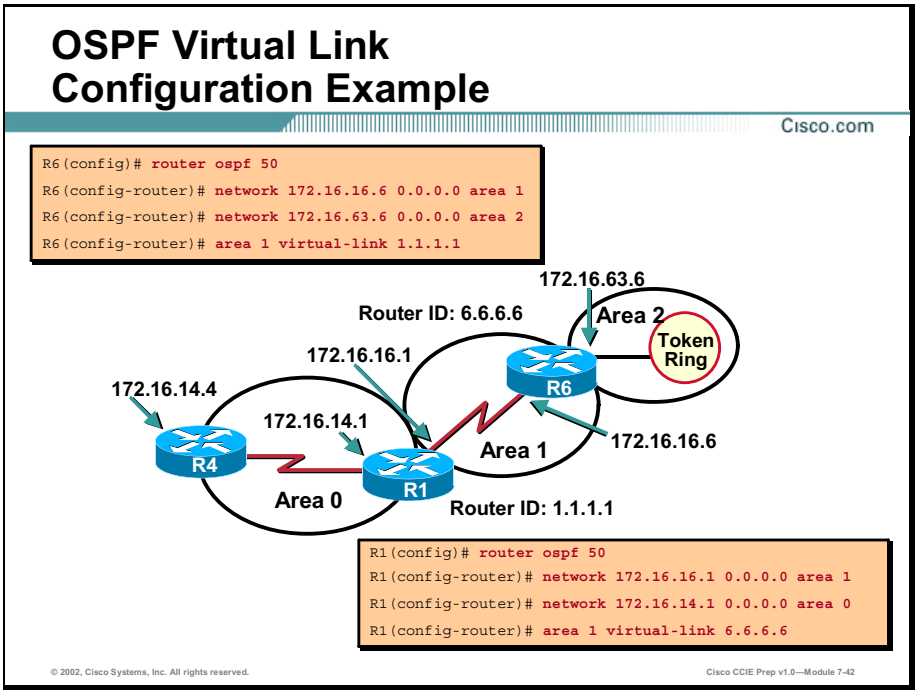

In this example, Area 2 does not have a direct physical connection to the backbone (Area 0). To provide connectivity to the backbone, a virtual link must be configured between R6 and R1. Area 1 will be the transit area, and R1 will be the entry point into Area 0. R6 will have a logical connection to the backbone through the transit area.

Both sides of the virtual link must be configured using the neighboring Area Boarder Router's (ABR's) router ID, not their physical interface's IP address. The neighboring ABR's router ID can be determined from the output of the **show ip ospf neighbor** command, or by telnetting to the neighboring ABR and using the **show ip ospf interface** command.

Table 6-18: < area <*area-id*> virtual-link <router-id>> Command

Command	Description
area <*area-id*> **virtual-link** <**router-id**>	Creates the virtual link by specifying the transit area and the router ID of the neighboring ABR. Performed in router configuration mode.

The need for a virtual link in your network can usually be determined by the following error:

```
Mar  1 07:02:19: %OSPF-4-ERRRCV: Received invalid packet: mismatch area ID, from
backbone area must be virtual-link but not found from 150.50.12.1,
Serial0/0
```

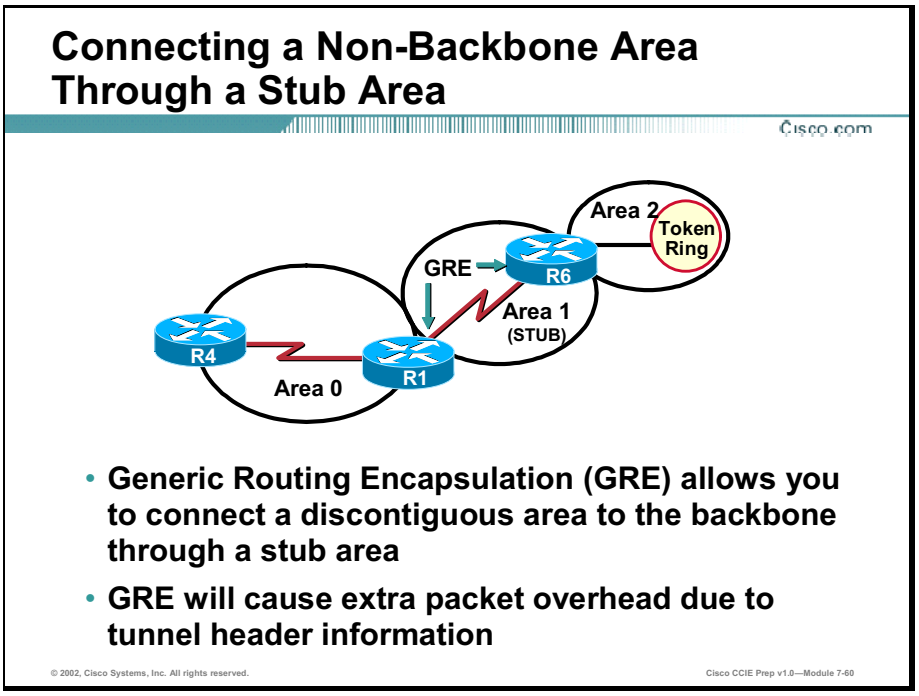

In this example, Area 1 has been configured as a stub area. This prevents the use of a virtual link, as virtual links are not allowed across stub areas. To provide Area 2 with connectivity to the backbone area, you could alternatively build a Generic Routing Encapsulation (GRE) tunnel between R6 and R1 and put the tunnel interfaces in area 0.

Table 6-19: GRE Tunnel Commands

Command	Description
interface tunnel *number*	Creates a virtual tunnel interface on the router.
tunnel source *interface | ip address*	Specifies the source of the point-to-point GRE tunnel. Can be specified by either the IP address or physical interface on the router.
tunnel destination *ip address*	Specifies the destination of the point-to-point GRE tunnel. This is the IP address the router on the other side of the tunnel.

The main differences between using a GRE tunnel or a virtual link to connect a discontiguous area to Area 0 are described in the following table:

Table 6-21: GRE Tunnel vs. Virtual Link

GRE Tunnel	Virtual Link
All traffic in the tunnel is encapsulated and decapsulated by the tunnel endpoints.	The routing updates are tunneled, but the data traffic is sent normally.
Tunnel headers in every packet cause overhead.	Data traffic is not subject to any tunnel overhead.
The tunnel can go through a stub area.	The transit area for a virtual link cannot be a stub area, because routers in the stub area will not have routes for external destinations.

OSPF Authentication

This section describes the use of plain text and Message Digest Version 5 (MD5) authentication to control neighbor adjacencies.

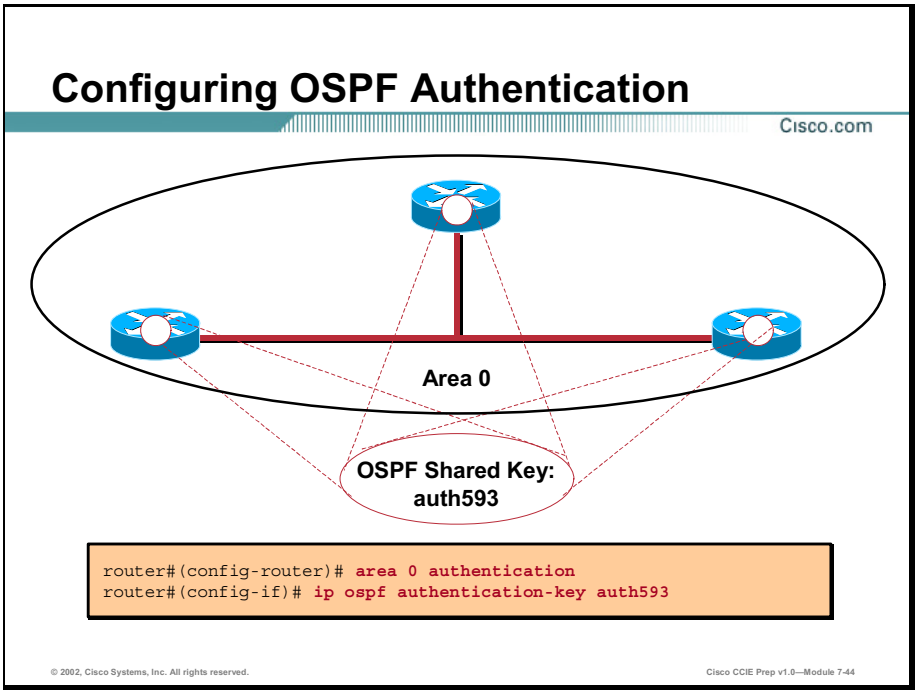

Plain Text Authentication

Plain text authentication allows a key (password) to be configured per area. All routers in the same area that want to participate in OSPF will have to be configured with the same key. Plain text authentication sends the authentication key itself in plain text over the wire. The drawback of this method is that it is vulnerable to eavesdropping attacks. Anybody with a protocol analyzer could easily get the plain text password off the wire.

Table 6-21

Command	Description
area *<area-id>* **authentication**	Enables plain text authentication for an OSPF area. When you configure authentication, you must configure all neighboring routers within the entire area for the same type of authentication.
ip ospf authentication-key *key*	Interface configuration command that defines the plain text key used between OSPF neighbors for authentication. Command applied in Global Interface mode.

Message Digest Authentication

MD5 authentication is a cryptographic form of authentication. A key (password) and key-id are configured on each neighboring router within the area. The router uses an algorithm based on the OSPF packet, the key, and the key-id to generate a "message digest hash" that gets appended to the packet. Since both neighbors are using the same key and key-id, they will be able to decode each other's hash. Unlike plain text authentication, the key itself is never exchanged over the wire. A non-decreasing sequence number is also included in each OSPF packet to protect against replay attacks.

Table 6-22

Command	Description
area <*area-id*> **authentication message-digest**	Enables MD5 authentication for an OSPF area. When you configure authentication, you must configure all neighboring routers within the entire area for the same type of authentication.
ip ospf message-digest-key *keyid* **md5** *key*	Interface configuration command that defines the key-id and key used to create the MD5 hash used between OSPF neighbors for authentication.

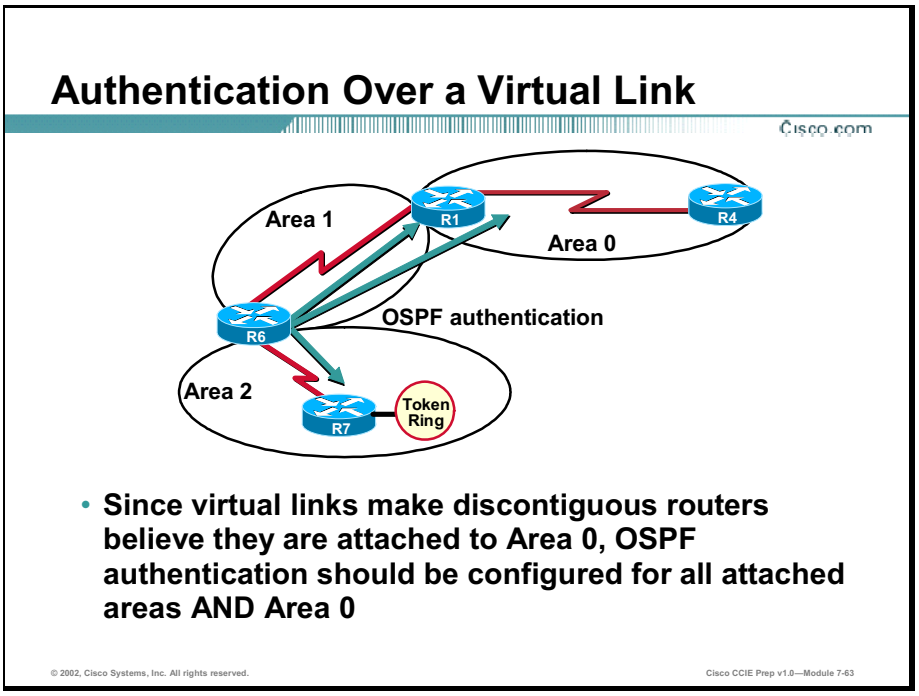

Virtual links support both plain text and MD5 authentication. There is one trick, however, to get authentication across a virtual link to work. When a virtual link is configured, the ABR of the area that is not physically connected to Area 0 now believes that it is part of area 0. Therefore, in addition to configuring authentication for the areas the ABR is attached to, authentication must also be configured for area 0 on the ABR.

Table 6-23

Command	Description
area *area-id* **virtual-link** *router-id* **authentication-key** *key*	Configures the plain text authentication key to be used across the virtual link
area *area-id* **virtual-link** *router-id* **message-digest-key** *key-id* **md5** *key*	Configures the key-id and key used to create the MD5 hash used to authenticate ABRs across a virtual-link

OSPF Demand Circuits

This section covers OSPF demand circuits.

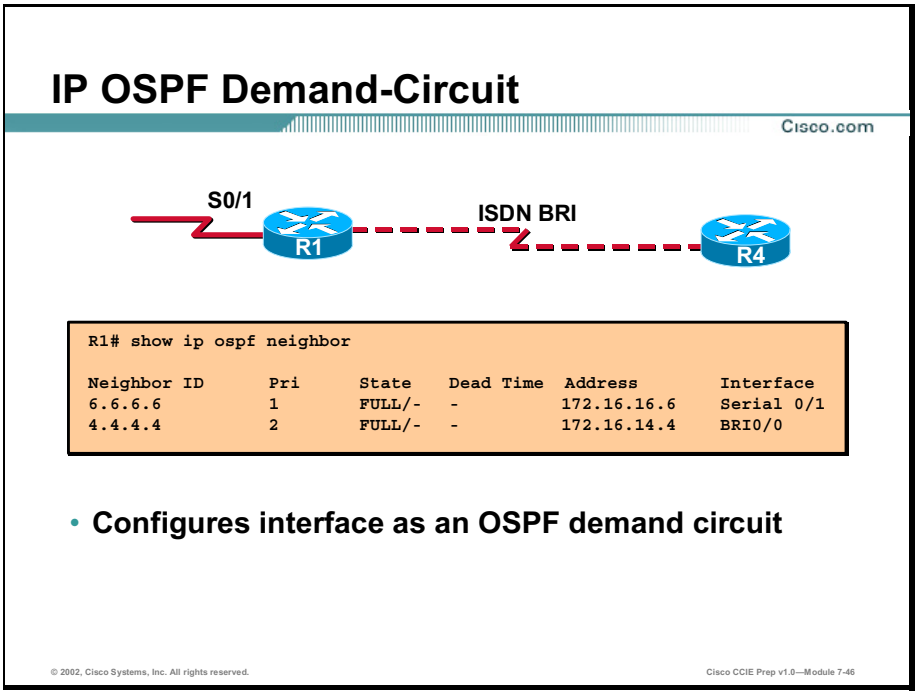

To create an OSPF demand circuit, only one side of the connection is required to have the **ip ospf demand-circuit** command under its interface. If the other side of the link is capable of understanding the DC bit, it automatically negotiates the demand circuit capability in the hello packets sent between the neighbors.

Table 6-24: < ip ospf demand-circuit > Command

Command	Description
`ip ospf demand-circuit`	Configures the interface to run as a demand circuit

After the first dead time interval (40 seconds by default on point-to-point links) the operation of the demand circuit can be verified with the **show ip ospf neighbor** command. If the demand circuit is functioning properly, you will not see a dead time being tracked for the neighbor on the other side of the demand circuit.

```
R1# show ip ospf neighbor

Neighbor ID     Pri   State     Dead Time   Address     Interface
6.6.6.6           1   FULL/  -       -      172.16.56.6  BRI0/0
4.4.4.4           1   FULL/  -    00:01:30  172.16.45.4  Serial0/0
```

Note	You can use the **ip ospf demand-circuit** command on any OSPF network type, however hellos are only suppressed on point-to-point or point-to-multipoint network types.

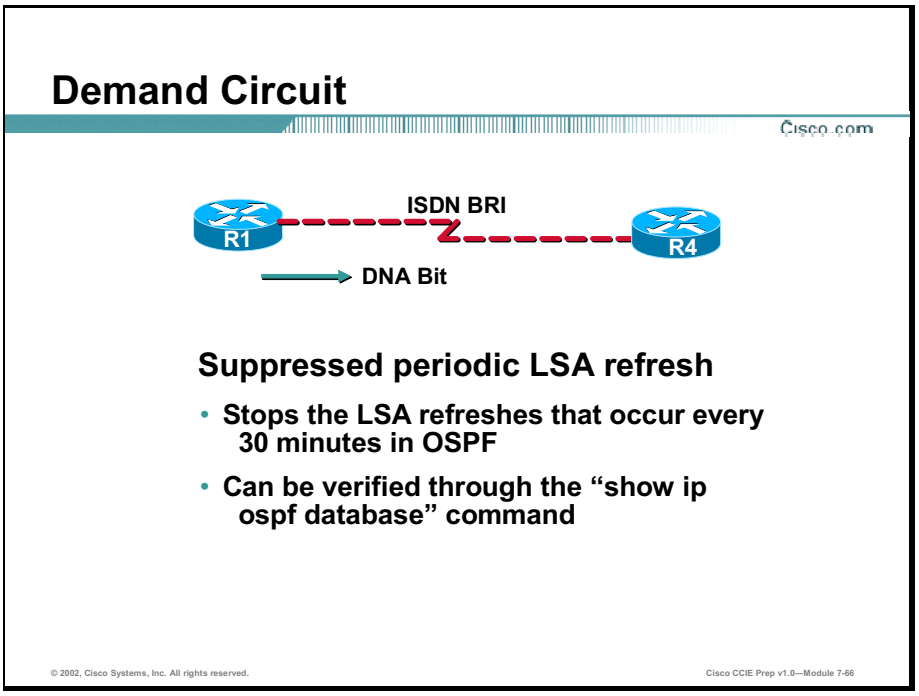

Suppressed Periodic LSA Refresh

The periodic Link-State Advertisement (LSA) refreshes that take place every 30 minutes in OSPF do not occur over the demand circuit. When the demand circuit is established, a unique option bit (the DC bit) is exchanged between the neighboring routers. If the two routers negotiate the Direct Current (DC) bit successfully, they will make a note of it and set a specific bit in the LSA Age field of LSAs they receive from the neighbor on the demand circuit. This specific bit is called the DoNotAge (DNA) bit. The DNA bit is the most significant bit in the LS Age field. By setting this bit, the LSA stops aging and no periodic updates are sent.

The setting of the DNA bit on LSAs can be verified by viewing the link-state database with the **show ip ospf database** command.

Summary

This section summarizes the key points discussed in this lesson.

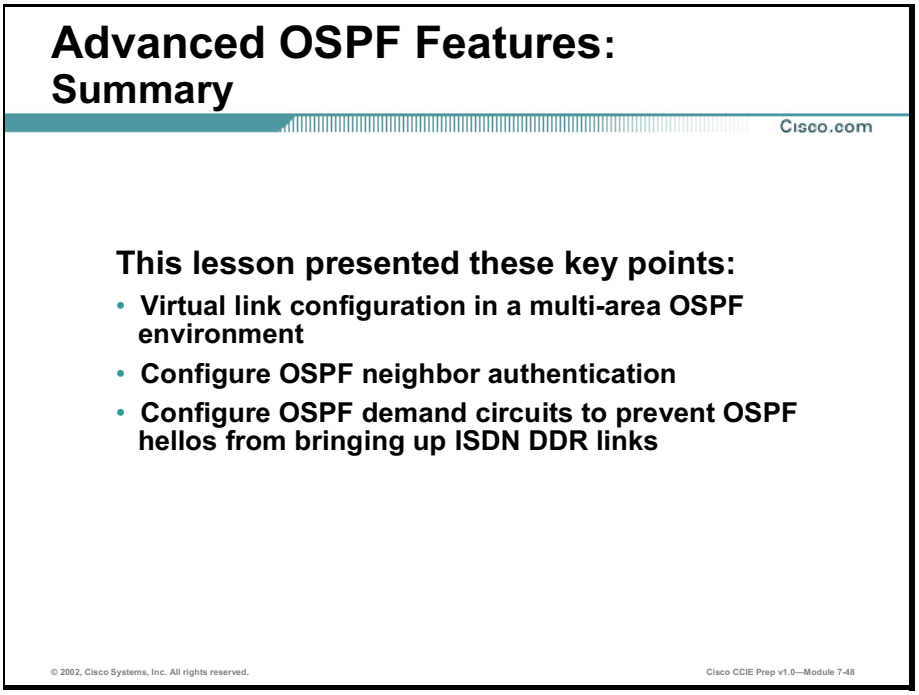

Next Steps

After completing this lesson, go to:

- Troubleshooting OSPF

References

For additional information, refer to these resources:

- http://www.cisco.com/warp/public/104/index.shtml - OSPF Technical Tips

Lesson Assessment (Quiz)

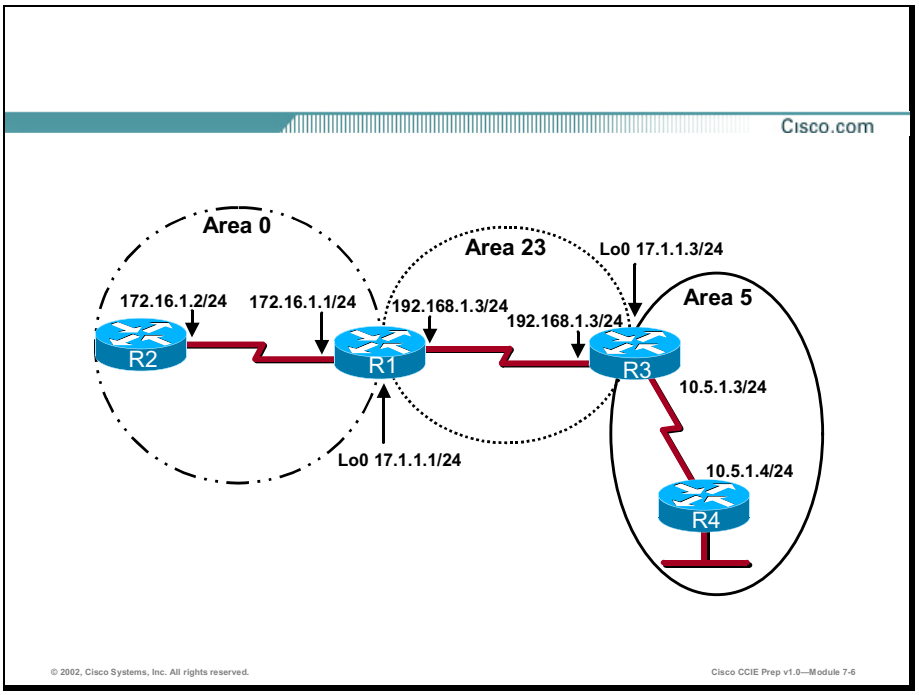

Q1) Based on the diagram above, what advanced OSPF feature is needed in this network?

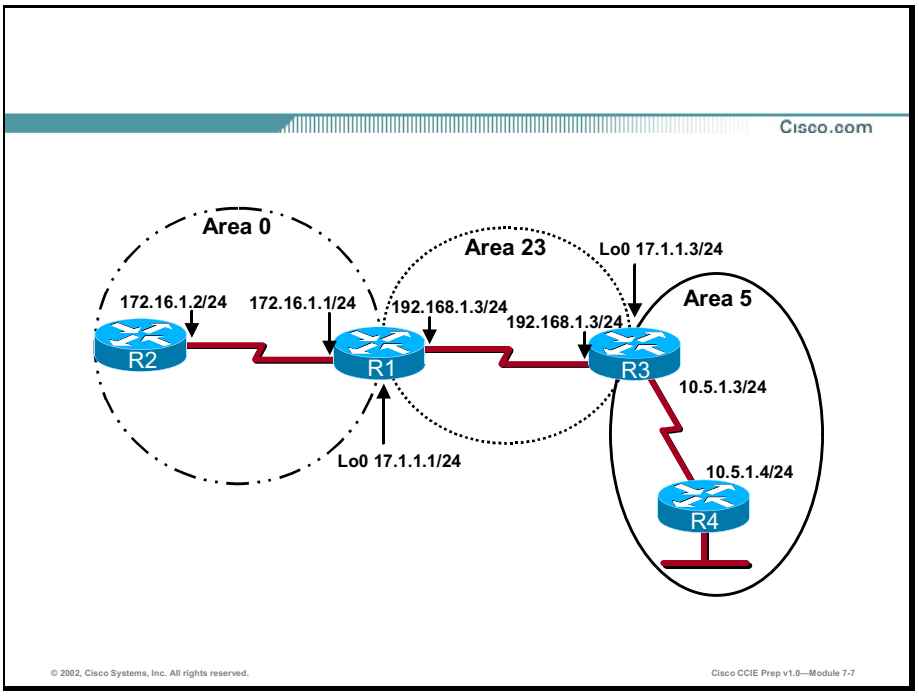

Q2) What is the correct command to create a virtual link on R1 in this diagram?

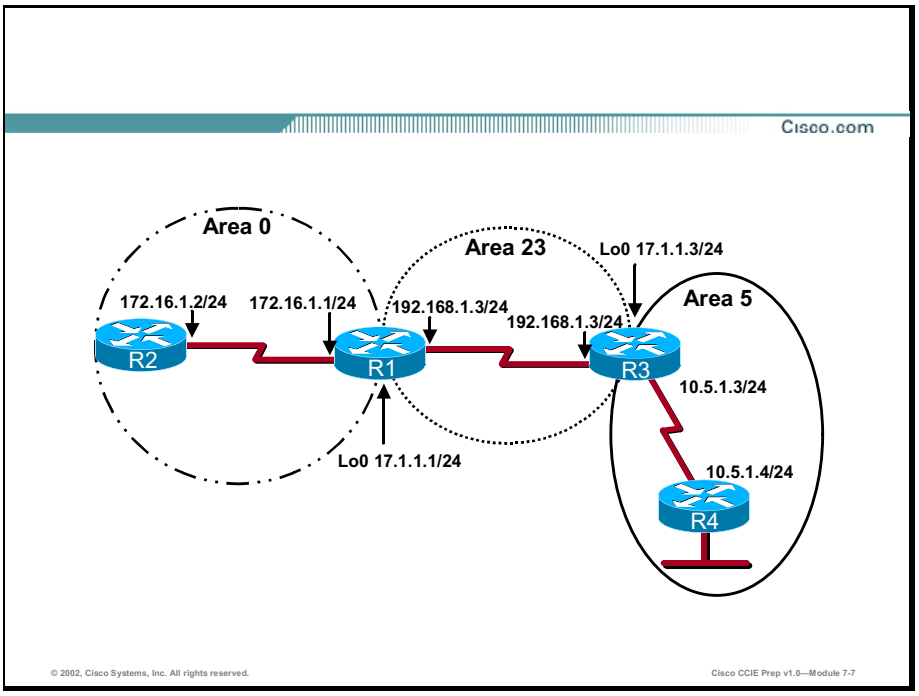

Q3) In what areas must authentication be configured for R4 in the above diagram?

Q4) LSAs that have been learned from a neighbor on an OSPF demand circuit are marked as what in the link-state database?

Troubleshooting OSPF

Overview

Because of the complexity of Open Shortest Path First (OSPF), Cisco provides many techniques to monitor and troubleshoot all areas of the protocol. This chapter discusses the basic troubleshooting commands as well as the advanced troubleshooting techniques that are necessary when working in the Cisco Certified Internetworking Expert (CCIE) lab environment.

Importance

OSPF is the core Interior Gateway Protocol (IGP) used in the CCIE lab.

Objectives

Upon completing this lesson, you will be able to:

- Verify OSPF neighbor adjacencies
- Verify the Link State Advertisements (LSAs) contained in the link-state database
- View the routes learned via OSPF
- Successfully troubleshoot a flapping OSPF demand circuit over Integrated Services Digital Network (ISDN)

Learner Skills and Knowledge

To fully benefit from this lesson, you must have these prerequisite skills and knowledge:

- Passed the CCIE Written qualification exam
- Completed the Building Scalable Cisco Internetworks (BSCI) course or have the equivalent knowledge

Outline

This lesson includes these sections:

- Overview
- Verifying OSPF Operation
- Troubleshooting a Flapping OSPF Demand Circuit over ISDN
- Summary
- Lesson Assessment (Quiz)

Verifying OSPF Operation

This section presents the commands used to verify OSPF operation.

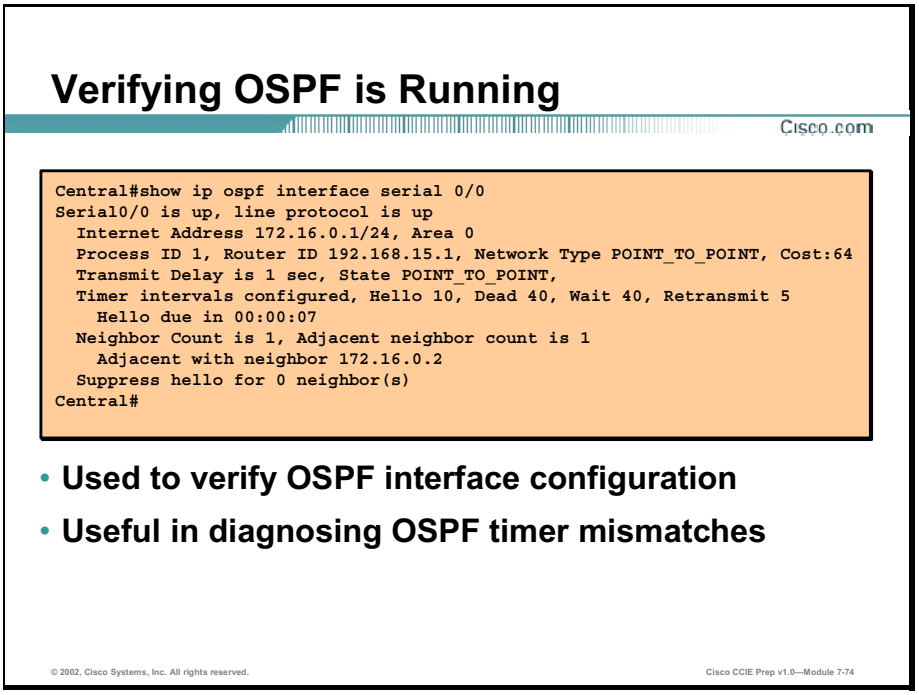

You can use the **show ip ospf interface** command to verify that Open Shortest Path First (OSPF) interfaces are running in the correct areas and have the correct OSPF network types defined. This command also displays other important information, such as the router ID, the cost of the interface, the Designated Router/Backup Designated Router (DR/BDR) of the segment (if applicable), the hello and dead timer intervals, if authentication is enabled for the interface, and the current neighbor adjacencies on the interface.

Table 6-25: < show ip ospf interface > Command

Command	Description
`show ip ospf interface`	Displays the interfaces on which OSPF is running and OSPF statistics about those interfaces.

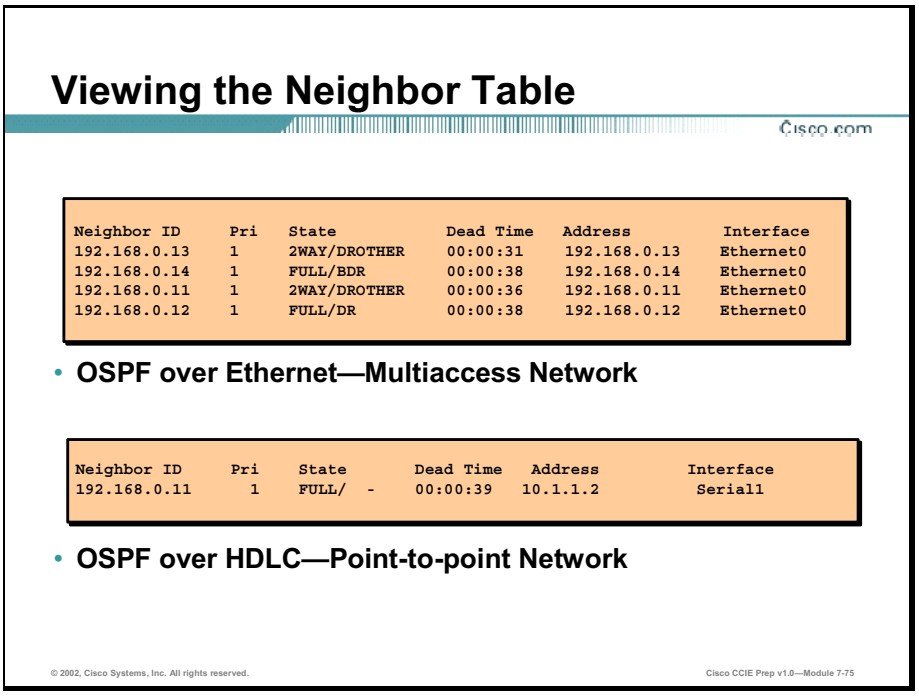

The **show ip ospf neighbor** command displays the OSPF neighbor database. This command can be used to verify the existence of OSPF neighbors including: their router IDs, their role on the segment (DR, BDR, or DROTHER), their current neighbor state (DOWN, ATTEMPT, INIT, 2WAY, EXSTART, EXCHANGE, LOADING, or FULL), and the interface off which they were learned.

Table 6-26: < show ip ospf neighbor [*type number*] [*neighbor-id*] **[detail]> Command**

Command	Description
show ip ospf neighbor *type number neighbor-id* **detail**	Displays the OSPF neighbor database. *type* – Optional keyword to display only neighbors off of a certain interface type. *number* – Optional keyword to display neighbors off of a certain interface. *neighbor-id* - Optional keyword to display only a certain neighbor. **detail** - Displays detailed information about neighbors.

The top output in the figure above is for OSPF on a broadcast multi-access Ethernet network. This is the expected output from a DROTHER (non DR/BDR) on the Ethernet segment. The neighbor states of FULL/DR and FULL/BDR indicate that this router has reached the FULL state with the DR and BDR. The lower output in the figure is for OSPF over a point-to-point network. A state of FULL/- indicates that this router has reached the full state with its neighbor and that there is no DR or BDR on this segment (because it is a point-to-point network).

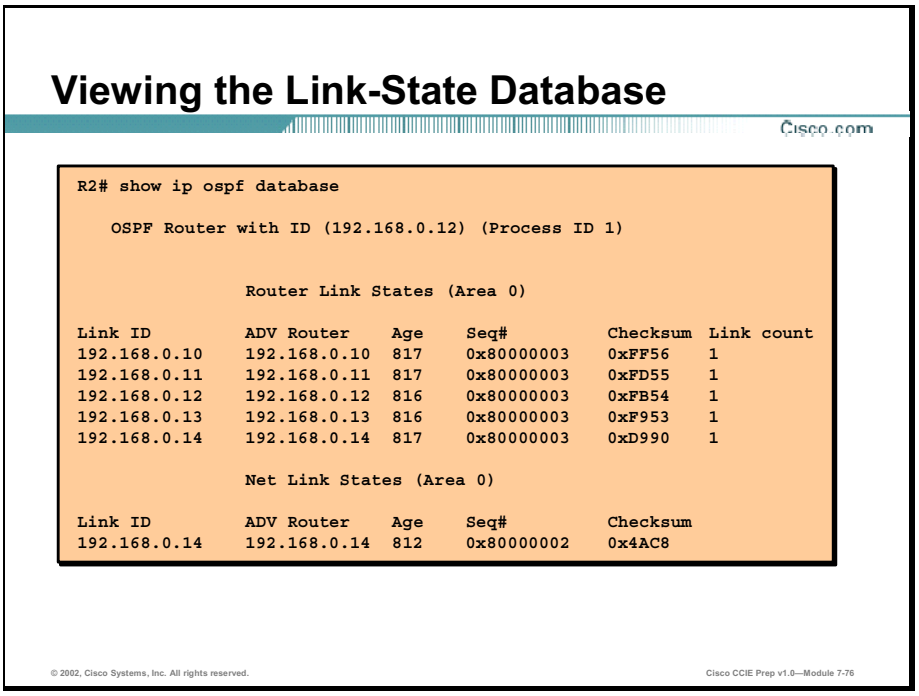

The **show ip ospf database** command displays the link-state database. The link-state database contains a listing of all the LSAs that a router knows about. This command is useful in verifying that OSPF is learning about a network, but is not putting it into the routing table for one reason or another. This command is also useful in verifying the operation of an OSPF demand circuit, by looking for LSAs marked as DoNotAge (DNA).

Table 6-27: < show ip ospf database > Command

Command	Description
`show ip ospf database`	Displays the link-state database on a router.

Note This command has many keywords that allow you to view only certain portions of the database by LSA Type. It is recommended that you become familiar with them to save time in troubleshooting OSPF problems.

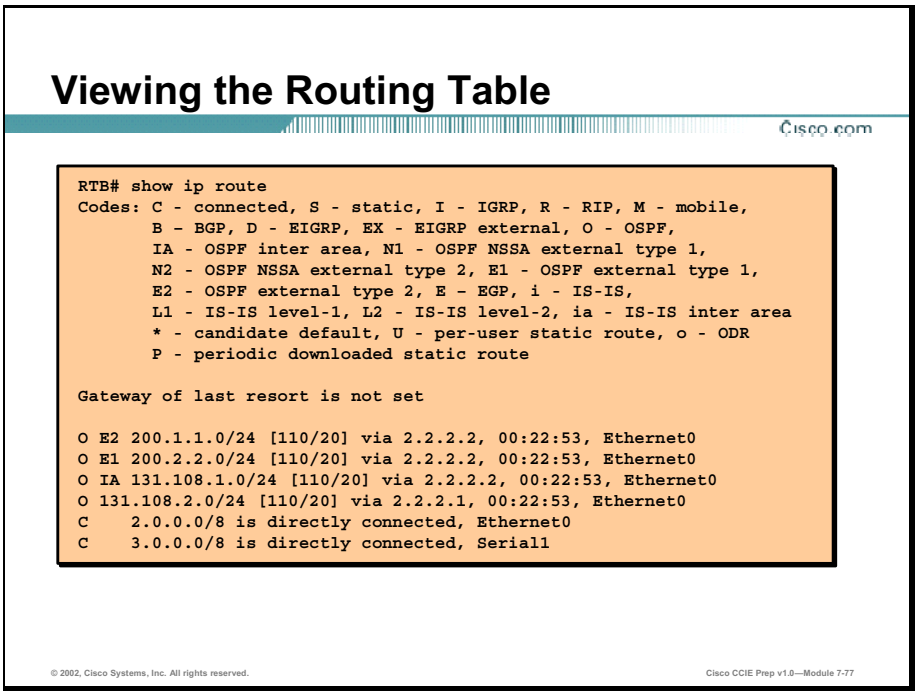

To view the routing table on a router, issue the **show ip route** command. To view only OSPF entries in the routing table, use the **show ip route ospf** command. OSPF entries in the routing table have several designations.

- O – (Intra-area OSPF route) Route to another network within the same area.
- IA – (Inter-area OSPF route) Route to a network in another area.
- E1 – (External Type 1 route) Route to a network that resides in another AS and was learned via redistribution. Type 1 routes have metrics that increment as they propagate throughout the OSPF domain.
- E2 – (External Type 2 route) Route to a network that resides in another AS and was learned via redistribution. Type 2 routes keep the metric assigned to them during redistribution as they propagate throughout the OSPF domain.

Table 6-28: < show ip route > Command

Command	Description
`show ip route`	Displays the routing table of a router.

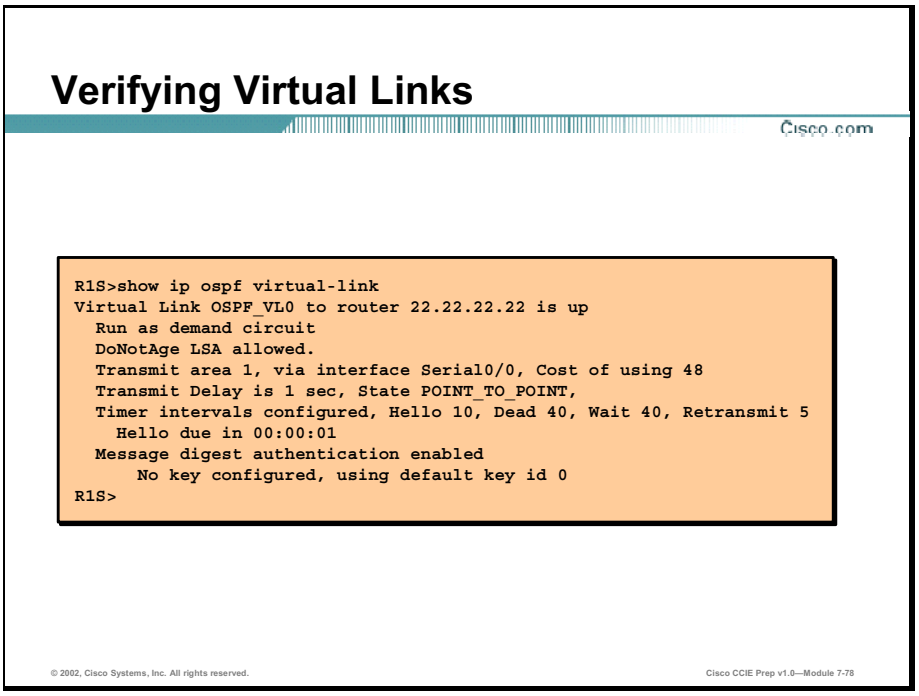

Use the **show ip ospf virtual-link** command to verify the status of a virtual link. This command will display the status of the virtual link, either up or down. This command also displays other important information, such as the Area Border Router (ABR) the virtual link is pointing to, the transit area the virtual link is running across, and information about authentication, if authentication has been configured on the virtual link.

Table 6-29: < show ip ospf virtual-link > Command

Command	Description
`show ip ospf virtual-link`	Displays the status of virtual links configured on the router.

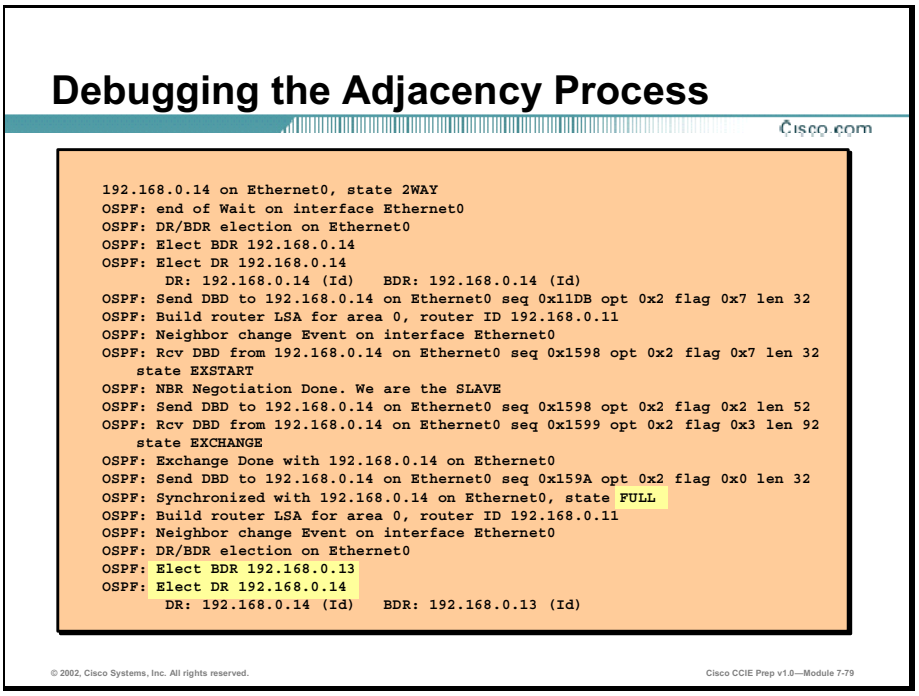

If an OSPF router is not forming a neighbor adjacency when it should, use the **debug ip ospf adj** command to troubleshoot the adjacency process. This command will display the neighbor adjacency states (DOWN, ATTEMPT, INIT, 2WAY, EXSTART, EXCHANGE, LOADING, and FULL) as they happen in real-time. This command usually tells you exactly why a neighbor adjacency is not being formed between two routers. Some of the most common reasons are: mismatched hello/dead timers, mismatched authentication parameters, or one router is configured with the stub flag and the other router is not.

Table 6-30: < debug ip ospf adj > Command

Command	Description
debug ip ospf adj	Debugs the OSPF neighbor adjacency process.

The **debug ip ospf adj** command is most helpful when an OSPF router first comes online. Since it is impossible to debug events on the router during boot up, use the **clear ip ospf process** command to manually restart the OSPF process on a router after issuing the **debug ip ospf adj** command.

Table 6-31: < clear ip ospf <*process-id*> **process > Command**

Command	Description
clear ip ospf <*process-id*> **process**	Manually clears the OSPF process on a router.

Troubleshooting a Flapping ISDN Link in OSPF

This section covers troubleshooting a flapping ISDN link in OSPF.

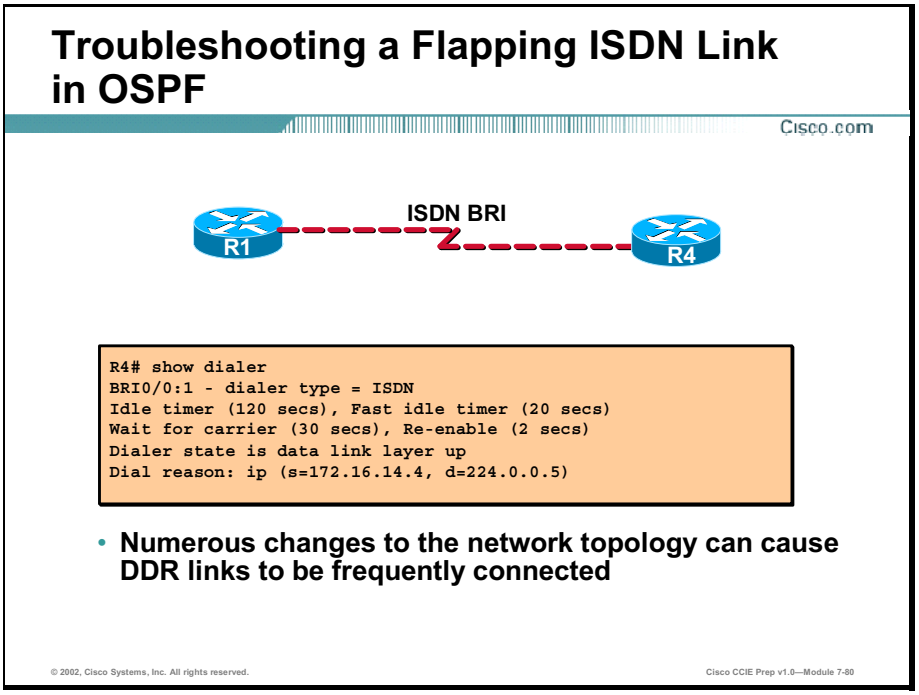

When an Integrated Services Digital Network (ISDN) link is configured as an OSPF demand circuit, OSPF hellos are suppressed and periodic LSA refreshes are not flooded over the link. These packets bring up the link only when they are exchanged for the first time, or when a change occurs in the network. The OSPF demand circuit feature allows the underlying Data-Link Layer to be closed when the OSPF network topology is stable. This is critical in a Dial-on-Demand Routing (DDR) environment.

In the diagram above, R4 and R1 are running an OSPF demand circuit across the ISDN link. The link between R4 and R1 keeps coming up, which defeats the purpose of the OSPF demand circuit feature. The output of the **show dialer** command shows that the link came up because of an OSPF hello packet.

```
R4# show dialer
BRI0/0:1 - dialer type = ISDN
Idle timer (120 secs), Fast idle timer (20 secs)
Wait for carrier (30 secs), Re-enable (2 secs)
Dialer state is data link layer up
Dial reason: ip (s=172.16.14.4, d=224.0.0.5)
```

OSPF could cause the ISDN link to activate for several reasons. These reasons will be discussed on the following pages.

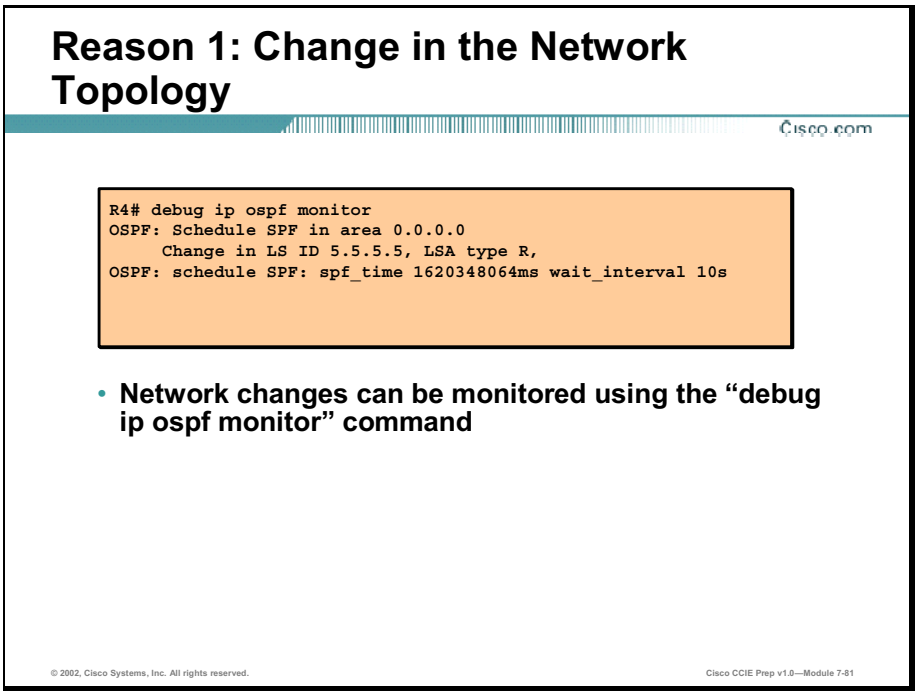

Reason 1: Change in the Network Topology

Whenever there is a change in an OSPF network topology, OSPF routers must be notified. In this situation, the OSPF demand circuit must be brought up, so that the neighbors can exchange the new LSA information. Once the new databases have been exchanged and synchronized, the link can go down again, and the adjacency remains in the FULL state.

Solution

To determine if the link is being brought up due to a change in network topology, use the **debug ip ospf monitor** command. It shows which LSA is changing, as shown below:

```
R4# debug ip ospf monitor
OSPF: Schedule SPF in area 0.0.0.0
     Change in LS ID 5.5.5.5, LSA type R,
OSPF: schedule SPF: spf_time 1620348064ms wait_interval 10s
```

The output above shows there was a change in the LSA with the router ID of 5.5.5.5, which causes the database to be resynchronized. If the network is stable, then this debug output will not display anything when the ISDN link comes up.

To reduce the chance of link flaps on the demand circuit, configure the area that contains the demand circuit as a stub or totally stubby area, if possible.

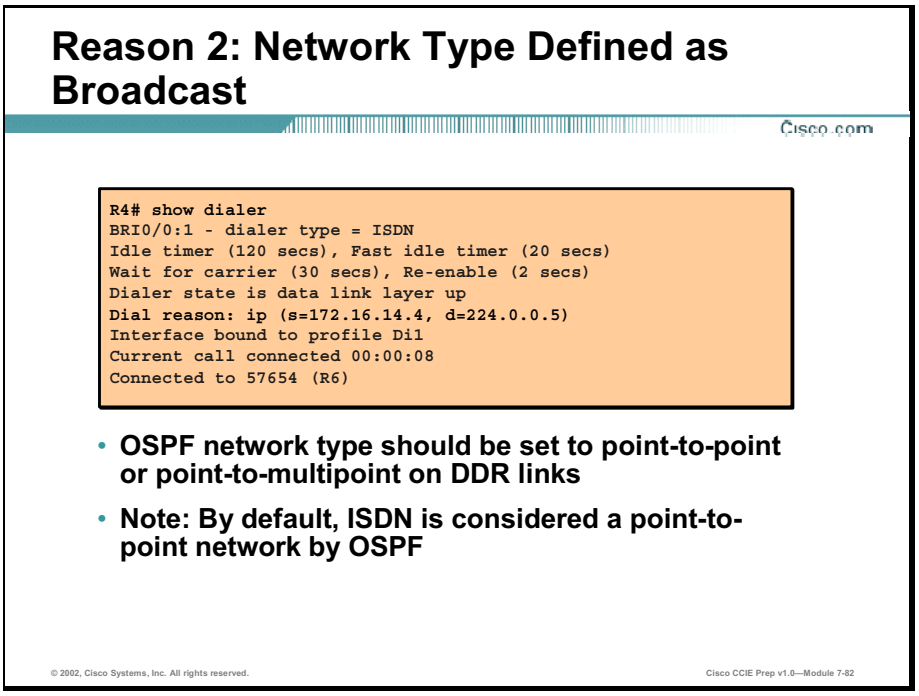

Reason 2: Network Type Defined as Broadcast

When you configure the OSPF demand circuit on a link, the OSPF network type must be defined as point-to-point or point-to-multipoint. Any other link type will cause the link to come up unnecessarily. This is due to the fact that OSPF hellos are not suppressed if the network type is anything other than point-to-point or point-to-multipoint. For example, with the network type defined as broadcast, OSPF hellos will bring the link up at every hello interval. The **show dialer** output here shows that the ISDN link was brought up because of an OSPF hello packet.

```
R4# show dialer
BRI0/0:1 - dialer type = ISDN
Idle timer (120 secs), Fast idle timer (20 secs)
Wait for carrier (30 secs), Re-enable (2 secs)
Dialer state is data link layer up
Dial reason: ip (s=172.16.14.4, d=224.0.0.5)
Interface bound to profile Di1
Current call connected 00:00:08
Connected to 57654 (R6)
```

Solution

To solve this problem, change the OSPF network type to either point-to-point or point-to-multipoint with the **ip ospf network** command. By changing the OSPF network type to point-to-point or point-to-multipoint, the OSPF hellos will be suppressed on the link, and the ISDN link will stop flapping.

Note By default, ISDN is considered a point-to-point network by OSPF.

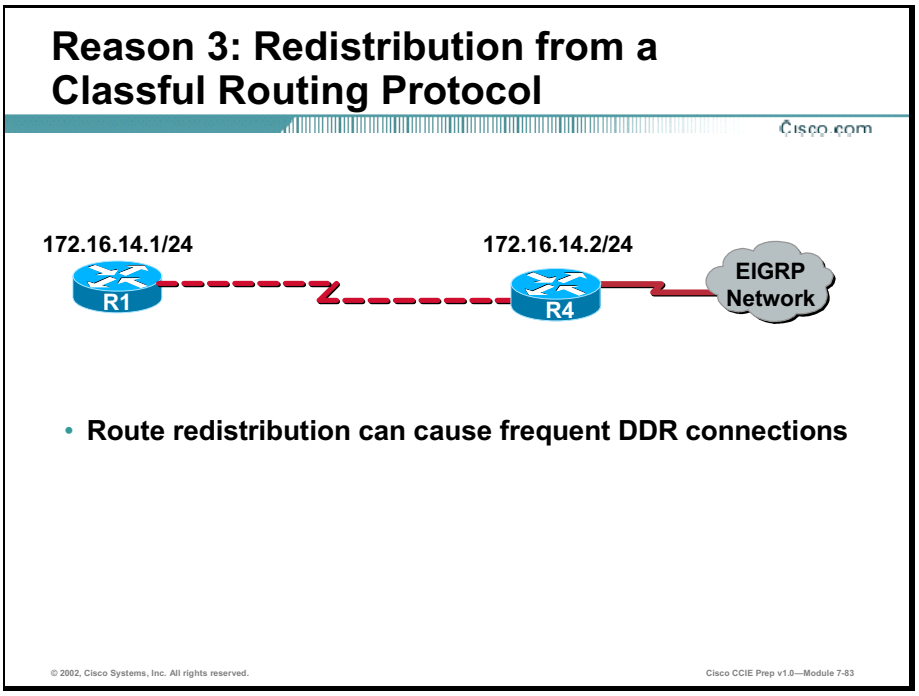

Reason 3: Redistribution from a classful routing protocol is being performed on one of the routers that connect to the OSPF demand circuit.

This is probably the most common reason that ISDN links configured as OSPF demand circuits still flap. It is also the hardest to troubleshoot and fix. In the example above, the ISDN link between R4 and R1 is 172.16.14.0/24 and is configured as an OSPF demand circuit. R4 is redistributing Enhanced Interior Gateway Routing Protocol (EIGRP) routes into OSPF. This causes many problems for an OSPF demand circuit on an ISDN link.

Since the encapsulation type on the ISDN link is Point-to-Point Protocol (PPP), both routers install a host route for the other side of the link as shown below.

```
R4# show ip route 172.16.14.1
Routing entry for 172.16.14.1/32
  Known via "connected", distance 0, metric 0 (connected, via interface)
  Routing Descriptor Blocks:
  * directly connected, via BRI0/0
      Route metric is 0, traffic share count is 1
```

EIGRP, IGRP, and Routing Information Protocol (RIP) are classful routing protocols. Therefore, the network statement in R4's EIGRP configuration is for the classful network of 172.16.0.0. This classful network statement causes the router to believe that the host route of 172.16.14.1/32 is being originated by EIGRP and gets redistributed into OSPF as an external route as shown below.

```
R4# show ip ospf database external 172.16.14.1
```

```
      OSPF Router with ID (4.4.4.4) (Process ID 1)

                Type-5 AS External Link States

LS age: 298
Options: (No TOS-capability, DC)
LS Type: AS External Link
Link State ID: 172.16.14.1 (External Network Number )
Advertising Router: 4.4.4.4
LS Seq Number: 80000001
Checksum: 0xDC2B
Length: 36
Network Mask: /32
       Metric Type: 2 (Larger than any link state path)
       TOS: 0
       Metric: 20
       Forward Address: 0.0.0.0
       External Route Tag: 0
```

Here is the problem: when the ISDN link goes down the /32 host route will disappear from the routing table. OSPF understands this as a change in topology and the demand circuit brings the link up again to propagate the MAXAGE version of the /32 host route to its neighbor. When the link comes up, the /32 mask gets inserted into the routing table again and the LSA age gets reset. After the first dead time interval on the link, the link goes down again. This process repeats itself, and the demand circuit link keeps flapping.

Solution: Use the no peer neighbor-route command

Under the ISDN BRI interface, configure the **no peer neighbor-route** command. This command prevents the /32 host route from being installed in the routing table. This command is only needed on the router performing redistribution, but is recommended for both sides of the ISDN link for consistency.

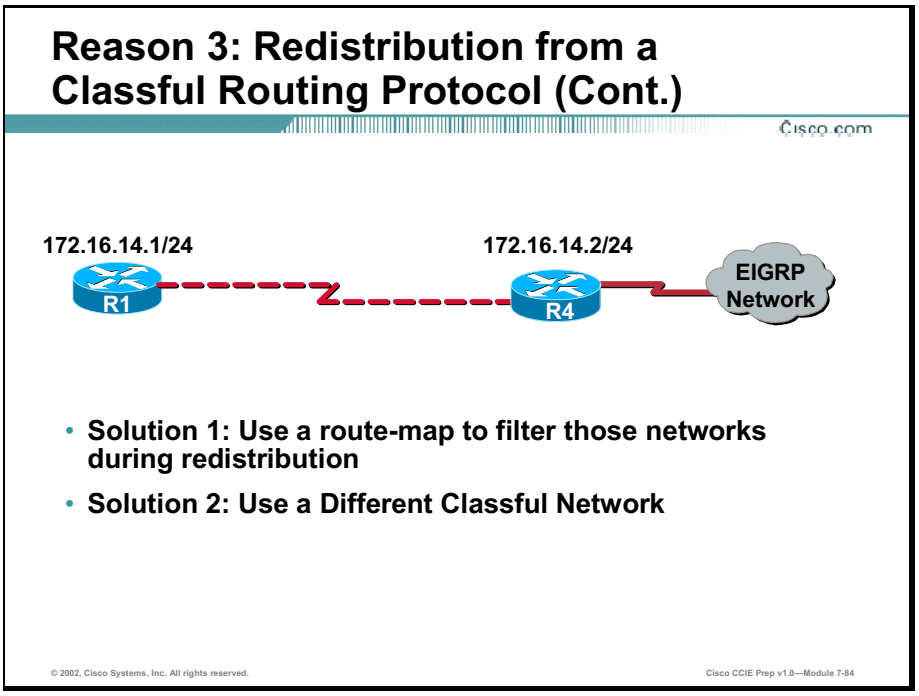

Reason 3 (cont.): Redistribution from a classful routing protocol is being performed on one of the routers that connect to the OSPF demand circuit

Other problems that will result from the redistribution of classful routing protocols into OSPF closely resemble the last problem. In the example above, the classful network statement of 172.16.0.0 in EIGRP on R4 also covers the serial interfaces on R4. Even though these networks live in OSPF, OSPF will think they are also external networks being redistributed from EIGRP. This will also cause the ISDN link to flap.

Solution 1: Use a route-map to filter those networks during redistribution

When redistributing from a classful protocol into OSPF, use a route-map to deny any networks that fall under the classful network space, but actually reside in OSPF.

First, we have to create an access list to match those networks. Then, we tie the access list to the route map using the **match address** parameter. Finally, we apply this route-map to the redistribution of EIGRP routes into OSPF on R4.

Note	The **match interface** parameter in the route-map can be used instead of the match address.

Solution 2: Use a Different Classful Network

Use a different classful network for the EIGRP domain. This will prevent all of the above problems.

Summary

This section summarizes the key points discussed in this lesson.

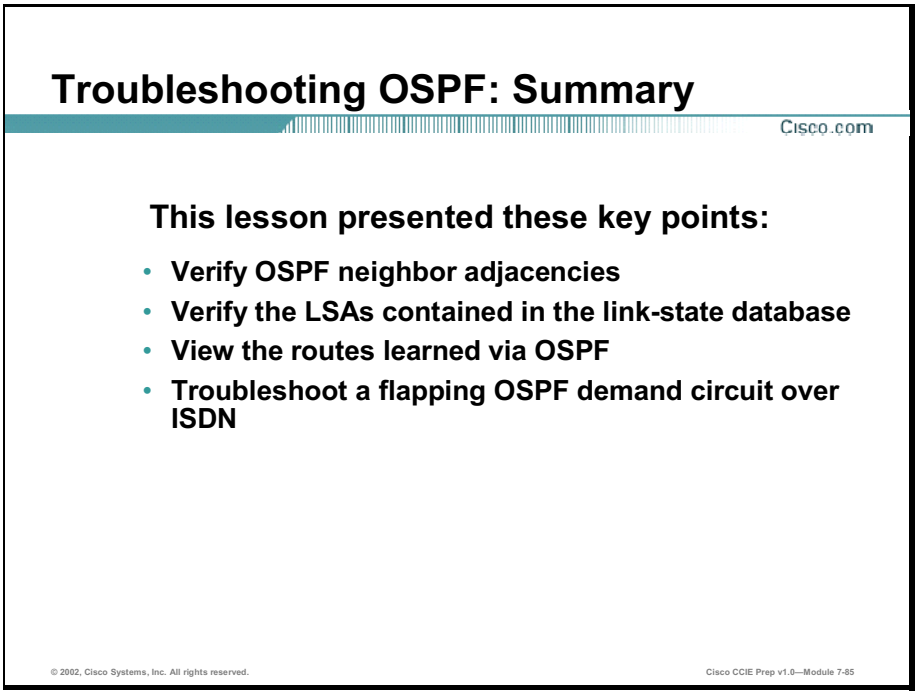

Next Steps

After completing this lesson, go to:

- Border Gateway Protocol (BGP) Technologies

References

For additional information, refer to these resources:

- http://www.cisco.com/warp/public/104/trouble_main.html

Lesson Assessment (Quiz)

Q1) What command is used to verify the area in which an interface belongs?

Q2) What command is used to view the OSPF neighbor table?

Q3) What command is used to view the router's link-state database?

Q4) What command is used to see OSPF neighbor adjacencies, as they are formed in real-time?

Q5) What command is used to verify if an OSPF demand circuit is being brought up due to a change in the link-state topology?

7

BGP Technologies

Overview

The Border Gateway Protocol (BGP) is a distance-vector routing protocol that you can use to exchange routing information among different Autonomous Systems (AS)s. This module examines the various topics and technologies used in a BGP environment.

Upon completing this module, you will be able to:

- Define BGP concepts and technologies
- Define internal BGP (iBGP)
- Define external BGP (eBGP)
- Define the different ways used to advertise networks in a BGP environment
- Configure the many advanced options of BGP
- Define the various **show** and **debug** commands used to troubleshoot BGP

Outline

The module contains these lessons:

- BGP Concepts
- eBGP Configuration
- Advertising Networks
- BGP Advanced Options
- Troubleshooting

iBGP Configuration

Overview

The Border Gateway Protocol (BGP) is a distance-vector routing protocol that is used to exchange routing information among different Autonomous Systems (AS). Understanding which BGP mode a router is using, is vital to proper configuration. This lesson will focus on internal BGP (iBGP) properties and proper configuration.

Importance

BGP version 4 is the current exterior routing protocol used on the Internet. When connecting a multi-homed enterprise to the Internet, it is vital that the administrator know BGP properties, configurations and troubleshooting tips.

Objectives

Upon completing this lesson, you will be able to:

- Describe a basic iBGP configuration
- Describe the iBGP “Rule of Synchronization”
- Describe the iBGP full mesh requirement
- Describe how route reflectors circumvent the full mesh requirement
- Use loopbacks for fault tolerance
- Understand when a BGP connection needs to be cleared

Learner Skills and Knowledge

To fully benefit from this lesson, you must have these prerequisite skills and knowledge:

- Passed the Cisco Certified Internetworking Expert (CCIE) written qualification exam
- Knowledge of Transfer Control Protocol/Internet Protocol (TCP/IP), access control lists, and Interior Routing Protocols (IGPs)

Outline

This lesson includes these sections:

- Overview
- BGP Functions
- Terminology
- BGP Path Selections
- Components
- iBGP Basic Configuration
- iBGP Advanced Configuration Rule of Synchronization
- Summary
- Lesson Assessment (Quiz)

BGP Functions

This section covers basic BGP functions.

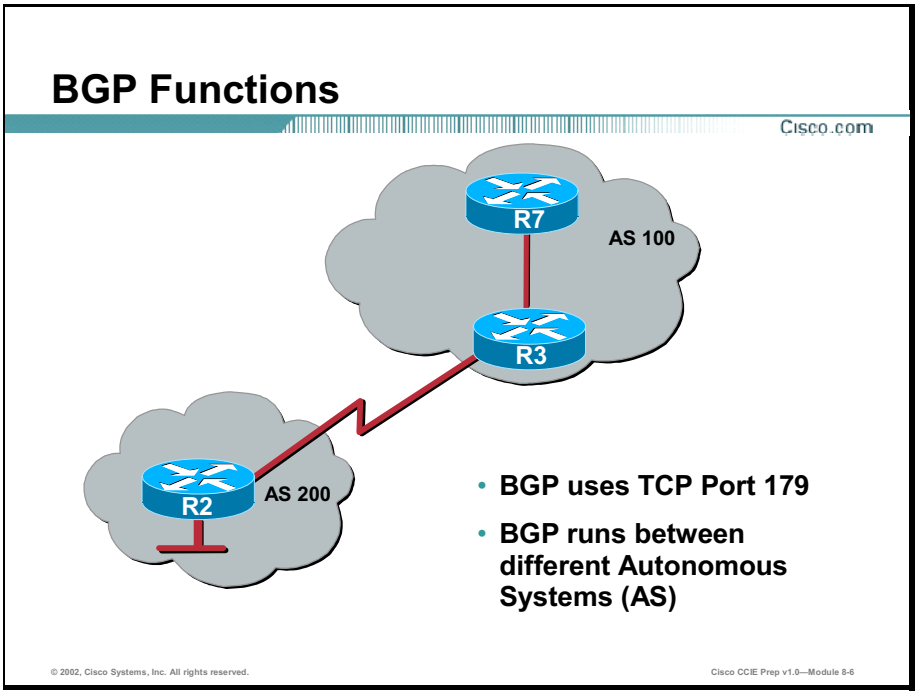

The Border Gateway Protocol (BGP) is an inter-autonomous system routing protocol. An Autonomous System (AS) is a network or group of networks under a common administration and with common routing policies. BGP is used to exchange routing information for the Internet and is the protocol used between Internet Service Providers (ISPs). Customer networks, such as universities and corporations, usually employ an Interior Gateway Protocol (IGP) such as Enhanced Interior Gateway Routing Protocol (EIGRP) or Open Shortest Path First (OSPF) for the exchange of routing information within their networks. Customers connect to ISPs, and ISPs use BGP to exchange customer and ISP routes.

The primary function of a BGP system is to exchange network reachability information with other BGP systems, including information about the list of autonomous system paths. This information can be used to construct a graph of autonomous system connectivity from which routing loops can be cut and with which autonomous system-level policy decisions can be enforced.

BGP neighbors exchange full routing information when the Transmission Control Protocol (TCP) (port 179) connection between neighbors is first established. When changes to the routing table are detected, the BGP routers send to their neighbors only those routes that have changed. BGP routers do not send periodic routing updates, and BGP routing updates advertise only the optimal path to a destination network.

Terminology

This section covers terminology associated with BGP.

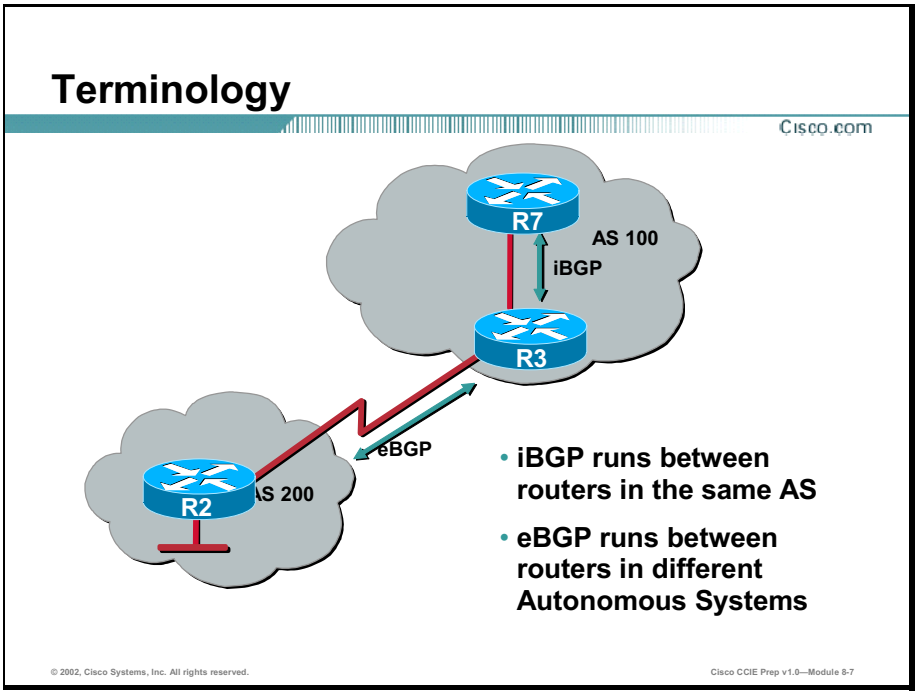

When BGP is used to exchange routing information between autonomous systems, the protocol is referred to as external BGP (eBGP). If BGP is used to exchange routes within an AS, then the protocol is referred to as interior BGP (iBGP). iBGP and eBGP will be discussed in the following lessons.

In BGP, each route consists of a network number, a list of autonomous systems that information has passed through (called the autonomous system path), and a list of other path attributes. Cisco supports BGP Versions 2, 3, and 4, as defined in Request for Comments (RFCs) 1163, 1267, and 1771, respectively.

BGP uses TCP as its transport protocol (specifically port 179). Any two routers that have opened a TCP connection to each other for the purpose of exchanging routing information are known as peers or neighbors.

iBGP is the form of BGP that exchanges BGP updates within an AS and eBGP is the form used to exchange updates when the BGP speakers are not in the same AS.

BGP Path Selection

This section will describe how BGP makes its selection for the best path.

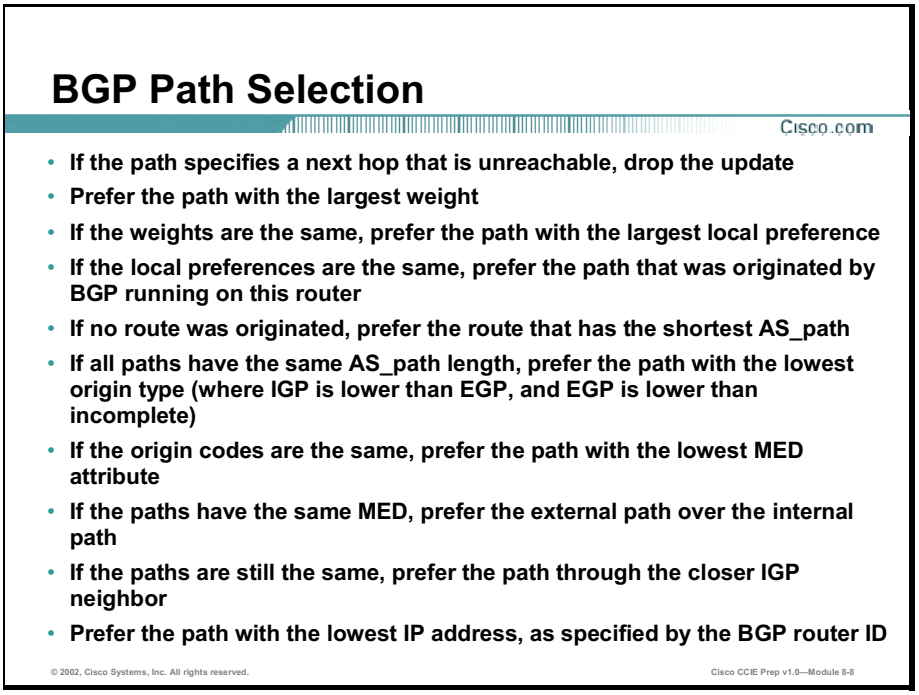

BGP could possibly receive multiple advertisements for the same route from multiple sources. BGP selects only one path as the best path. When the path is selected, BGP puts the selected path in the Internet Protocol (IP) routing table and propagates the path to its neighbors. BGP uses the criteria, in the order presented, to select a path for a destination. The following pages will cover the more common components used in BGP path selection.

Components

This section covers the components used with BGP, which affect path selection.

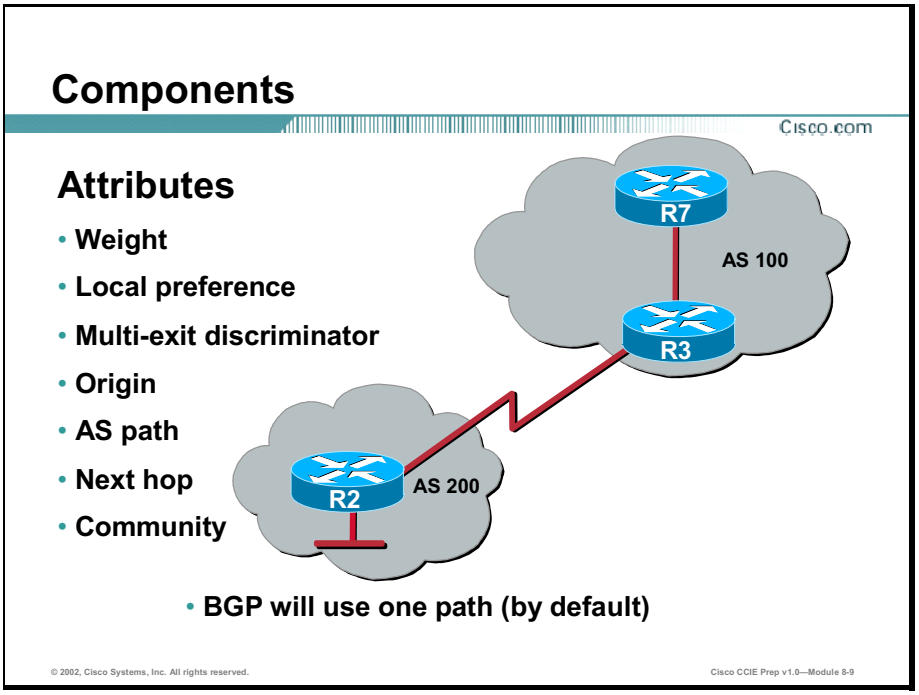

When a BGP speaker learns two identical eBGP paths for a prefix from a neighboring AS, it will choose the path with the lowest route-id as the best path. This best path is installed in the IP routing table. If BGP multi-path support is enabled and the eBGP paths are learned from the same neighboring AS, multiple paths are installed in the IP routing table instead of picking one best path.

During packet switching, either per-packet or per-destination load balancing is performed among the multiple paths, depending on the switching mode. A maximum of six paths is supported. The **maximum-paths** router configuration command controls the number of paths allowed. By default, BGP will install only one path to the IP routing table.

Attributes

Routes learned via BGP have associated properties that are used to determine the best route to a destination when multiple paths exist to a particular destination. These properties are referred to as BGP attributes. An understanding of how BGP attributes influence route selection is required for the design of robust networks.

iBGP Basic Configuration

This section covers basic iBGP configuration.

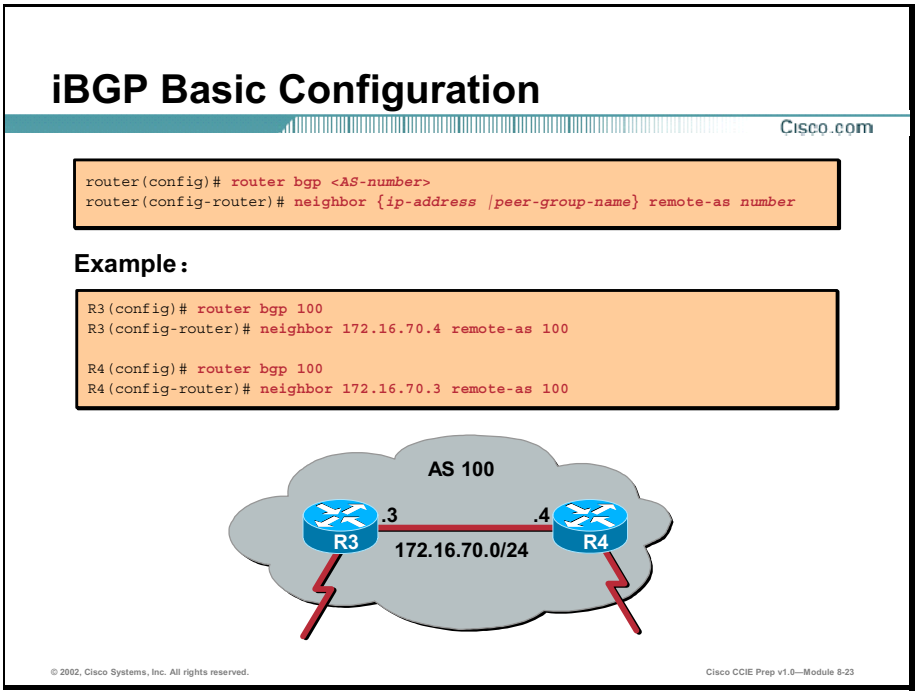

There are two primary commands required to configure an iBGP neighbor relationship.

router bgp *<AS-number>*

neighbor *{ip-address | peer-group-name}* **remote-as** *AS-number*

The first command **router bgp** *<AS-number>* is used to enable the router as a BGP speaker and place the router in an autonomous system.

The second command is used to create an iBGP (or eBGP) TCP neighbor relationship with another router. You cannot exchange routing updates without an established neighbor relationship. iBGP is used between routers in the same AS. An iBGP neighbor relationship can occur between routers that are not directly connected. This is a very important concept as most students new to BGP think of neighbors as directly connected. Remember, all you need to establish a neighbor relationship is TCP connectivity, because distance is not an obstacle.

Once the TCP connection is up, the routers send open messages in order to exchange values such as the AS number, the BGP version they are running, the BGP router ID and the keepalive hold time. After these values are confirmed and accepted, the neighbor connection is established. Any state other than "established" is an indication that the two routers did not become neighbors, and BGP updates will not be exchanged.

When identifying your neighbor, you can do so in two ways:

- Via any active IP address assigned to an interface

- Via a loopback address (preferred)

In the previous scenario, an IP address of the peer is used for creating an iBGP neighbor relationship. The purpose is to form an iBGP neighbor relationship between routers R3 and R4.

The important items to remember are each router is in AS100 and is identifying its peer as being in AS100. This configuration creates an iBGP neighbor relationship. Notice the physical Ethernet interface of the peers is being used, although any active IP address on each router would suffice. The closest IP address is used for simplicity.

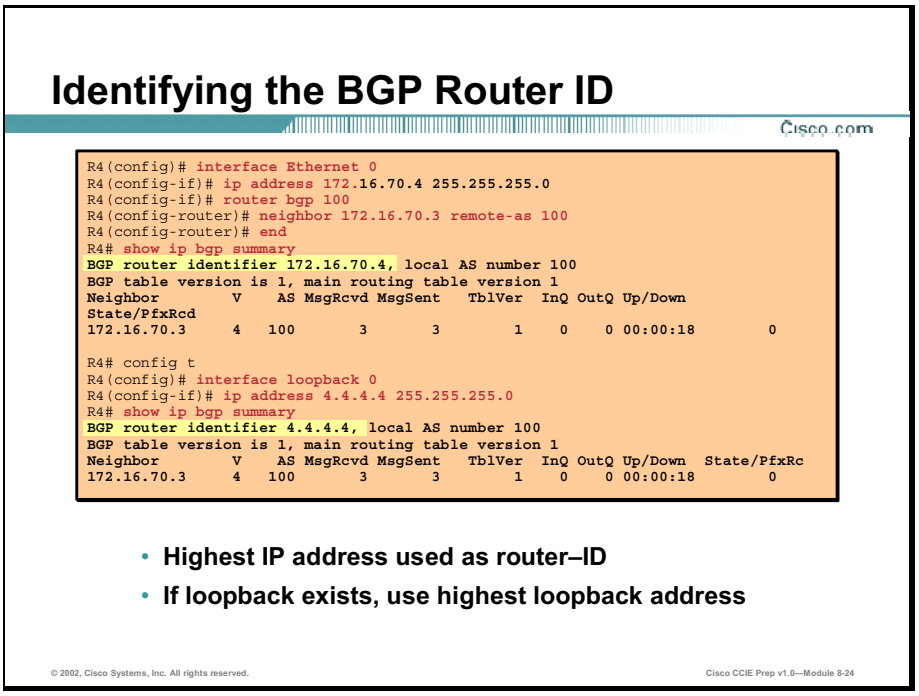

When you start BGP on a router using the **router bgp <as-number>** command, the router automatically assigns the physical interface with the highest IP address as the router ID. If loopbacks are configured, then the loopback with the highest IP address is used instead. This is demonstrated in the example.

```
R3(config)# interface Ethernet 0
R3(config-if)# ip address 172.16.70.3 255.255.255.0
R3(config-if)# router bgp 100
R3(config-router)# neighbor 172.16.70.4 remote-as 100

R4(config)# interface Ethernet 0
R4(config-if)# ip address 172.16.70.4 255.255.255.0
R4(config-if)# interface loopback 0
R4(config-if)# ip address 4.4.4.4 255.255.255.0
R4(config-if)# router bgp 100
R4(config-router)# neighbor 172.16.70.3 remote-as 100
```

Now, to discover the Router ID's, perform a **show ip bgp summary** on each router.

```
R3# show ip bgp summary
BGP router identifier 172.16.70.3, local AS number 100
BGP table version is 1, main routing table version 1
Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.70.4     4   100       4       4        1    0    0 00:01:31        0
```

Notice the BGP router identifier is 172.16.70.3 for R3, the IP address of the Ethernet 0 interface.

```
R4# show ip bgp summary
BGP router identifier 4.4.4.4, local AS number 100
BGP table version is 1, main routing table version 1
Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRc
172.16.70.3     4   100       3       3        1    0    0 00:00:18        0
```

R4 has both a physical interface as well as a loopback interface. Notice, the loopback interface has been chosen as the BGP router identifier for R4.

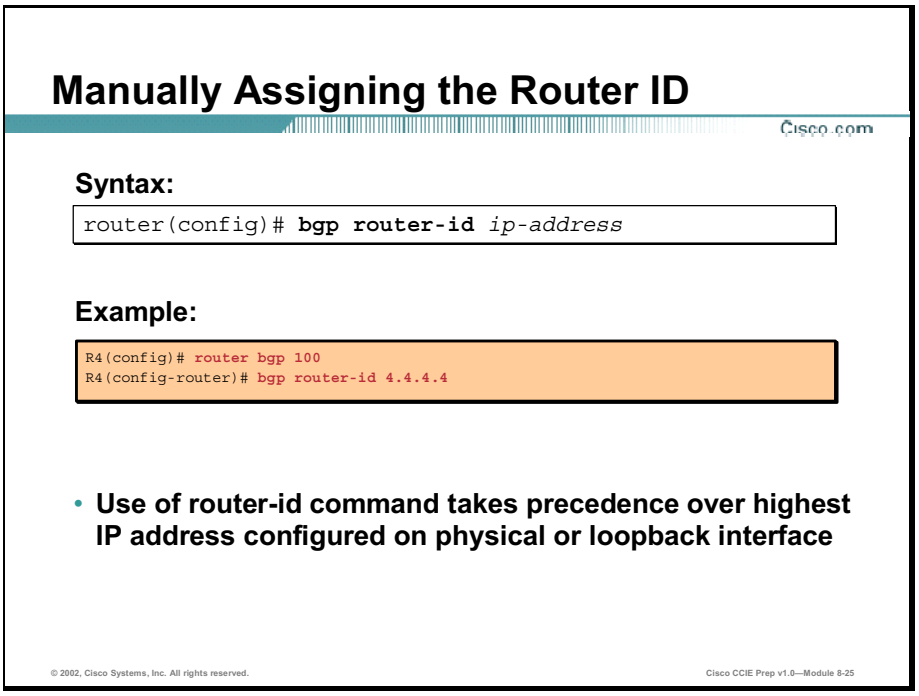

When the IP address of a physical interface is used as the router ID, it can pose a problem. If the physical interface should go down BGP neighbor relationships can be lost. It is possible to change the router ID manually or use a loopback interface. Doing so prevents instability in BGP connections due to interface flapping, but can also cause another problem when used with OSPF. BGP and OSPF will perform redistribution if they agree on the same router ID. If you manually change the router ID on BGP, OSPF will need to have the same router ID if you wish redistribution to occur.

The syntax to "hard code" the router ID is the following:

bgp router-id *ip-address*

This enables you to have a router ID that will never change. Look at R4's configuration again.

```
R4(config)# interface Ethernet 0
R4(config-if)# ip address 172.16.70.4 255.255.255.0
R4(config-if)# interface loopback 0
R4(config-if)# ip address 4.4.4.4 255.255.255.0
R4(config-if)# router bgp 100
R4(config-router)# neighbor 172.16.70.3 remote-as 100
```

At this time, the BGP router identifier is 4.4.4.4, the IP address of the loopback interface. If you never wanted this to change, you could do the following:

```
R4(config)# router bgp 100
R4(config-router)# bgp router-id 4.4.4.4
```

Then, if you add a loopback address, such as:

```
R4(config)# interface loopback 5
R4(config-if)# ip address 144.144.144.144 255.255.255.255
```

Normally, when you reload the router, the higher loopback address of 144.144.144.144 would become the Router ID, but due to the command **bgp router-id 4.4.4.4**, the Router ID will never change.

iBGP Advanced Configuration Rule of Synchronization

When an AS provides transit service to other ASs and there are non-BGP routers in the AS, transit traffic might be dropped, if the intermediate non-BGP routers have not learned routes for that traffic via an IGP.

This leads to one of the important rules of BGP. It is called the BGP "Rule of Synchronization."

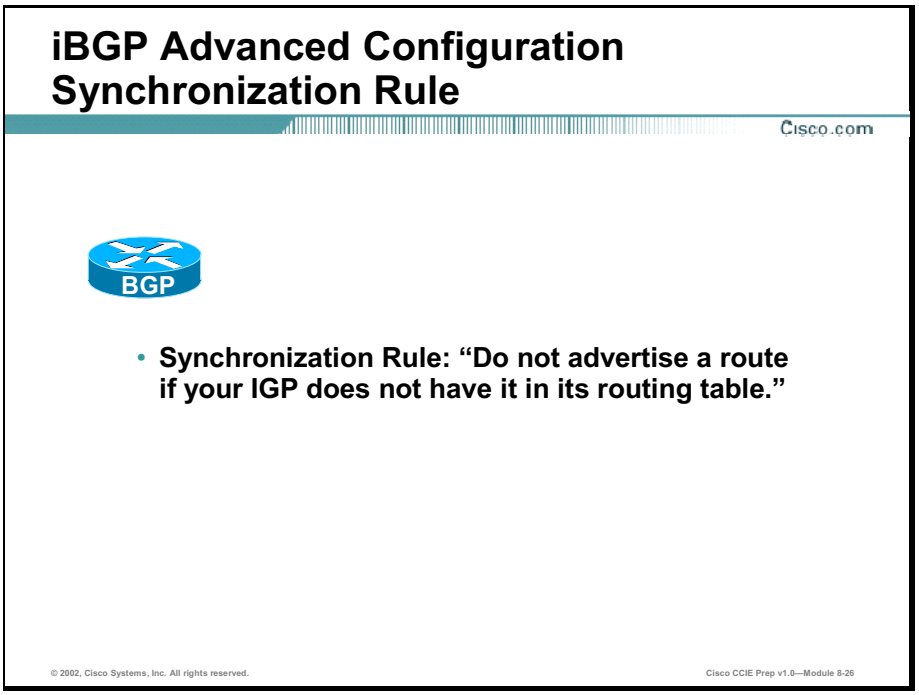

The BGP rule of synchronization states that if an AS provides transit service to another AS, BGP should not advertise a route until all of the routers within the AS have learned about the route via an IGP. In other words, it states "Do not advertise a route if the IGP does not have it in its routing table."

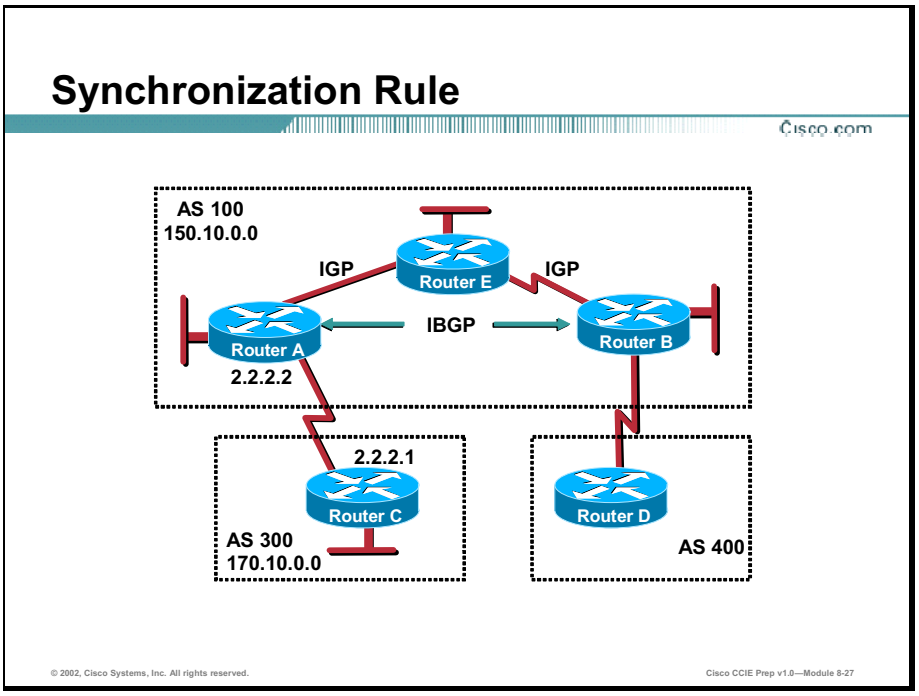

This is an example of the rule of synchronization and how it works. In the diagram, Router C sends updates about network 170.10.0.0 to Router A. Routers A and B are running iBGP, so Router B receives updates about network 170.10.0.0 via iBGP. If Router B wants to reach network 170.10.0.0, it sends traffic to Router E. If Router A does not redistribute network 170.10.0.0 into an IGP, Router E has no way of knowing that network 170.10.0.0 exists and will drop the packets.

If Router B advertises to AS 400 that it can reach 170.10.0.0 before Router E learns about the network via IGP, traffic coming from Router D to Router B with a destination of 170.10.0.0 will flow to Router E and be dropped.

RFC1771 Introduction

"To characterize the set of policy decisions that can be enforced using BGP, one must focus on the rule that a BGP speaker advertise to its peers (other BGP speakers which it communicates with) in neighboring ASs only those routes that it itself uses."

This situation is handled by the synchronization rule of BGP, which states that if an AS (such as AS 100 in the diagram) passes traffic from one AS to another AS, BGP does not advertise a route before all routers within the AS (in this case, AS 100) have learned about the route via an IGP. In this case, Router B waits to hear about network 170.10.0.0 via an IGP before it sends an update to Router D. In some cases, you might want to disable synchronization. Disabling synchronization allows BGP to converge more quickly, but it might result in dropped transit packets.

You can disable synchronization if one of the following conditions is true:

- Your AS does not pass traffic from one AS to another AS.

- All the transit routers in your AS run BGP.

For example, to turn off synchronization on Router A, issue the commands:

```
RouterA(config)router bgp 100
RouterA(config-router)no synchronization
```

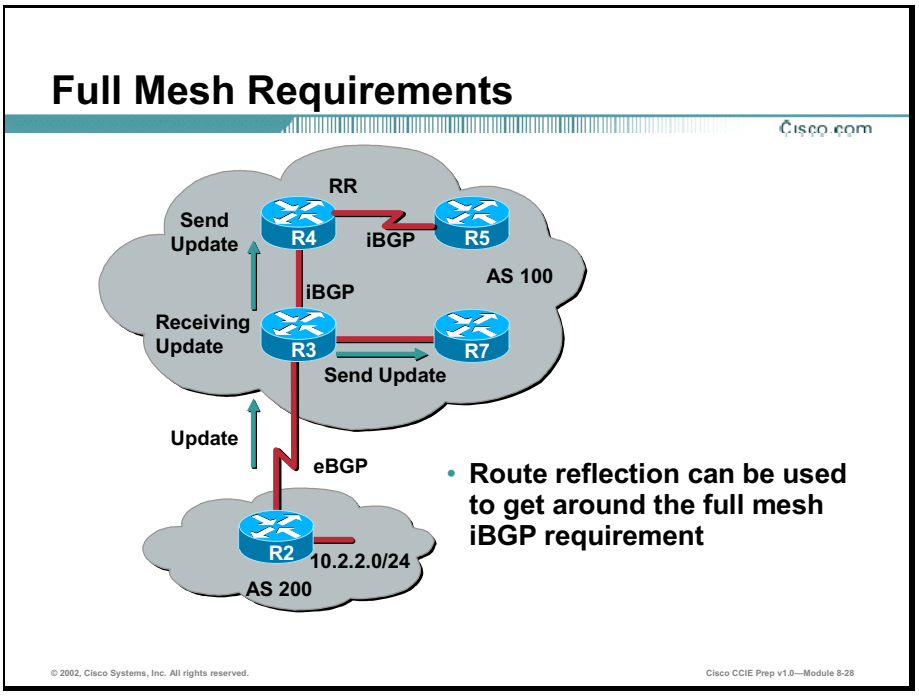

Once a BGP session has been established, updates are exchanged to provide all the locally known routes with only the best path advertised. Incremental update messages are exchanged later. If the best path is received from an eBGP peer, then it is advertised to all peers. This is another very important concept to understand. If a BGP speaker receives an update from an eBGP peer, it will send that update which is the same as one-hop into the iBGP domain.

In the diagram, R2 is part of AS200 and is connected to R3, which is part of AS100. R3 also has iBGP peering relationships with R4 and R7, but not with R5. R4 has an iBGP peer relationship with R5.

Now, see what happens when an update for network 10.2.2.0/24 is sent to R3.

R3 receives the update and sends it to its iBGP peers, which are R4 and R7. R4, since it received the update from its iBGP peer R3, will not send the update to R5. Consequently, R5 will never learn about the 10.2.2.0/24 network.

RFC1771 Section 9.2.1 Internal Updates

"When a BGP speaker receives an UPDATE message from another BGP speaker located in its own autonomous system, the receiving BGP speaker shall not re-distribute the routing information contained in that UPDATE message to other BGP speakers located in its own autonomous system."

This leads to the requirement that in iBGP, you must have a full mesh. R3 will peer with R5, which means R5 would have received the update. If a full mesh is not created, missing routes and other troubles can develop. The requirement of a full mesh for a small iBGP environment is not a problem, but when you are dealing with an iBGP domain that uses hundreds of routers, it becomes a huge scaling problem. Having each router peer with ever other router can cause

serious problems with the routers, namely in memory and Central Processing Unit (CPU) cycles. For example, in an AS with 100 BGP speakers, they will be required to build 4950 iBGP sessions. The calculation used is: N=devices; N(N-1)/2.

When receiving updates and the best path comes from an iBGP peer, it should be advertised only to eBGP peers. Again, a full iBGP mesh should be created.

As described earlier, a BGP speaker does not advertise a route learned from another iBGP speaker to a third iBGP speaker. Route reflectors ease this limitation and allow a router to advertise (reflect) iBGP-learned routes to other iBGP speakers, thereby reducing the number of iBGP peers within an AS.

neighbor *{ip-address | peer-group-name}* **route-reflector-client**

Look at the configuration, but this time configure R4 to be a route reflector for the client R5.

On R4, add the following configuration command:

```
R4(config)# router bgp 100
R4(config-router)# neighbor 172.16.45.5 route-reflector-client
```

Notice, you configured only the route reflector server (R4) not the client (R5). Now, when R4 receives the update from R3, it will “reflect” the update to its client R5.

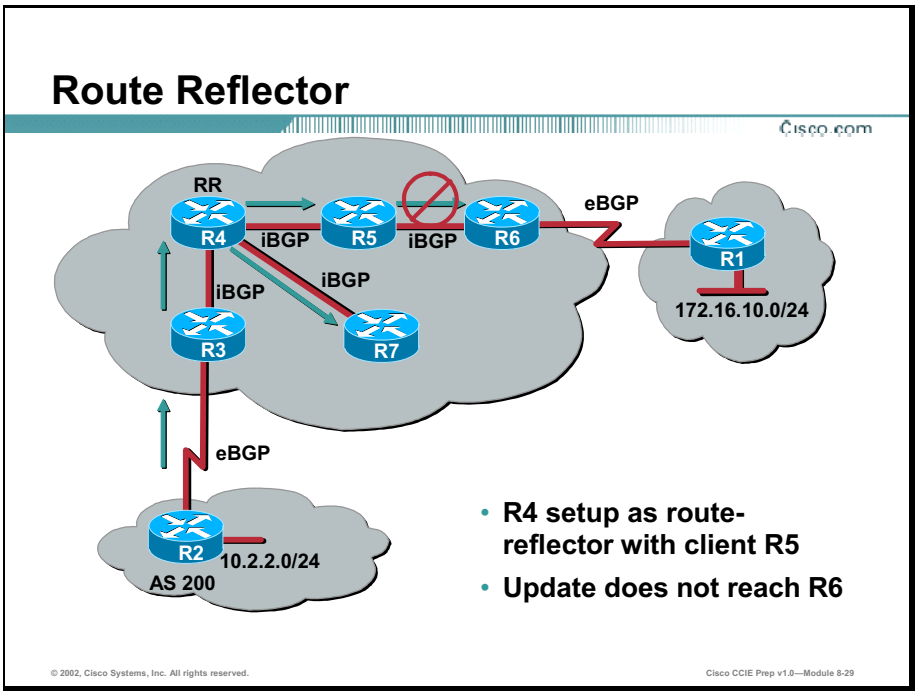

When you configure a route reflector client, you must also remember that the route reflector server will send updates to its clients *as well* as non-clients. Non-clients are iBGP peers without a specific route-reflector-client configuration statement. To understand this concept, take a look at a more complex example. Here R3, R4, R5, R6, and R7 are part of AS100. First, configure R4 as a route reflector with R5 as its client.

The example shows what happens when R2 sends an update for network 10.2.2.0/24 to R3. R3 will receive its update and send an update to its iBGP neighbors. In this case, R3 only has one neighbor (R4). R4 receives the update and because it is configured as a route reflector server, sends its update to R5, the client. R4 will also send the update to its non-clients, which in this case is R7. So with R4 configured as a route reflector server, R5 and R7 will receive the update to network 10.2.2.0/24.

But, there is a problem. R5 will not send the update it receives to R6, and what happens when R1 sends its update to R6 for network 172.16.10.0/24?

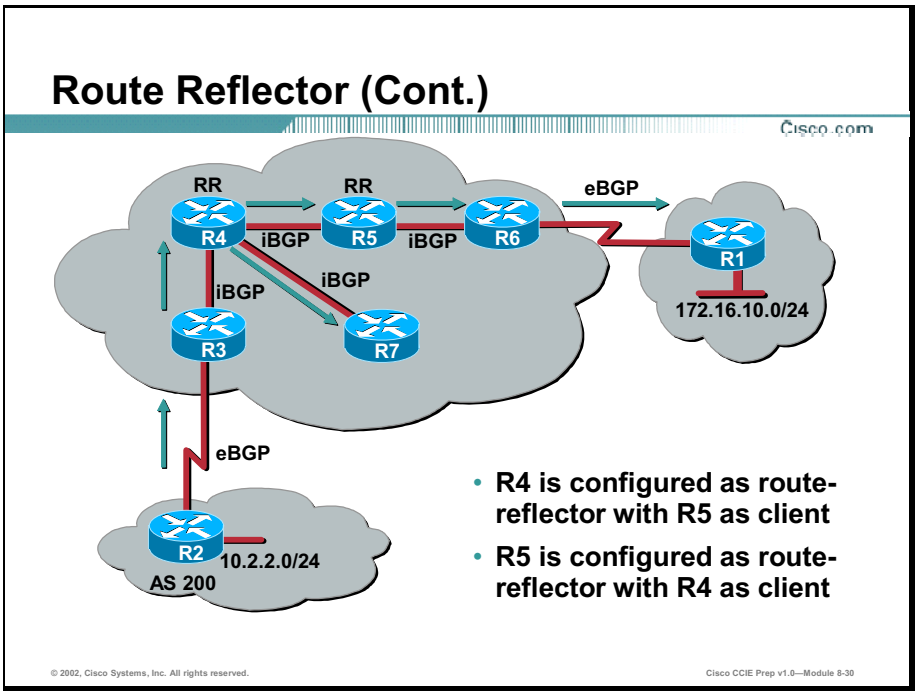

The solution to the full mesh problem might not seem obvious at first, but what happens when R5 is configured as a route reflector server with R4 as its client?

The example shows what happens when R2 sends its update for network 10.2.2.0/24 to R3.

R3 receives the update and sends the update to its iBGP peer R4.

R4 is configured as a route reflector server with R5 as client, sends the update to R5 and R7 as its non-client.

Since R5 is also configured as a route reflector server with R4 as its client, it will receive the update and send an update to its client R4 and its non-client, which in this case is R6.

R6 receives the update and sends an update to its eBGP neighbor R1.

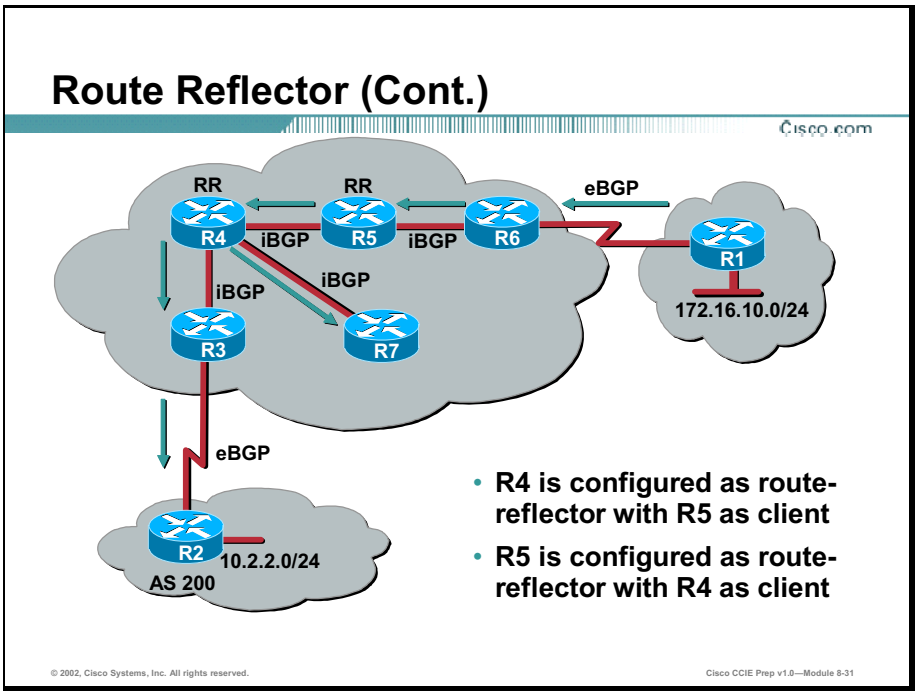

But, there is still a problem. What happens to the updates coming from R1? Trace an update of network 172.16.10.0/24 coming from R1 going to R6.

R6 receives the update and sends the update to its iBGP peer R5.

R5 is configured as a route reflector server with R4 as its client and sends the update to R4.

R4 receives the update and because it is configured as a route reflector server with R5 as its client, it will send an update to its client (R5), as well as its non-clients (R3 and R7).

R3 receives the update and sends an update to is eBGP neighbor R2.

Now, you have simulated a complete iBGP mesh and can receive updates from the eBGP neighbors, learn in the AS, and send them to the other eBGP neighbor.

As you can see, strategically placing route reflectors in your environment can overcome the full mesh requirement of iBGP.

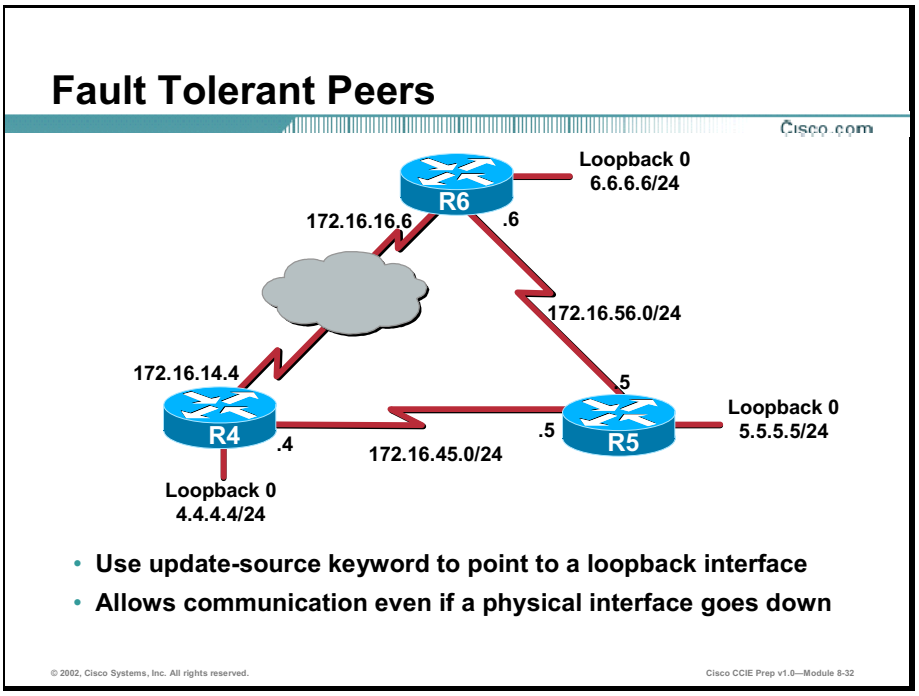

neighbor *{ip-address | peer-group-name}* **update-source** *interface-name*

Using the **update-source** keyword for a peer allows iBGP sessions to use any operational interface for TCP connections. iBGP neighbor relationships can occur as long as there is TCP connection between peers. Using physical interfaces can create problems when they go down and another link is active to the peer. In other words, there is no fault-tolerance. Using loopback interfaces can allow fault tolerance in the BGP domain even when a link fails.

Using a loopback interface to define neighbors is common with iBGP, but not with eBGP. Normally the loopback interface is used to make sure the IP address of the neighbor stays up and is independent of properly functioning hardware. In the case of eBGP, the peer routers are directly connected frequently, and loopback does not apply.

The example shows a configuration where using loopback interfaces will allow fault tolerance.

In this scenario, R4, R5, and R6 are running a common IGP and have full routing tables for all routes including the loopbacks. They also are running iBGP in a full mesh.

```
R4(config)# int loopback 0
R4(config-if)# ip address 4.4.4.4 255.255.255.0
R4(config-if)# router bgp 100
R4(config-router)# neighbor 5.5.5.5 remote-as 100
R4(config-router)# neighbor 5.5.5.5 update-source loopback 0
R4(config-router)# neighbor 6.6.6.6 remote-as 100
R4(config-router)# neighbor 6.6.6.6 update-source loopback 0

R5(config)# int loopback 0
R5(config-if)# ip address 5.5.5.5 255.255.255.0
```

```
R5(config-if)# router bgp 100
R5(config-router)# neighbor 4.4.4.4 remote-as 100
R5(config-router)# neighbor 4.4.4.4 update-source loopback 0
R5(config-router)# neighbor 6.6.6.6 remote-as 100
R5(config-router)# neighbor 6.6.6.6 update-source loopback 0

R6(config)# int loopback 0
R6(config-if)# ip address 6.6.6.6 255.255.255.0
R6(config-if)# router bgp 100
R6(config-router)# neighbor 4.4.4.4 remote-as 100
R6(config-router)# neighbor 4.4.4.4 update-source loopback 0
R6(config-router)# neighbor 5.5.5.5 remote-as 100
R6(config-router)# neighbor 5.5.5.5 update-source loopback 0
```

Look at the configuration leading to the use of R4's loopback address as the peer IP address. Notice that R5 and R6 are referring to R4's loopback address in their neighbor statements. R4 is telling both R5 and R6 that it is using its loopback 0 as its update source.

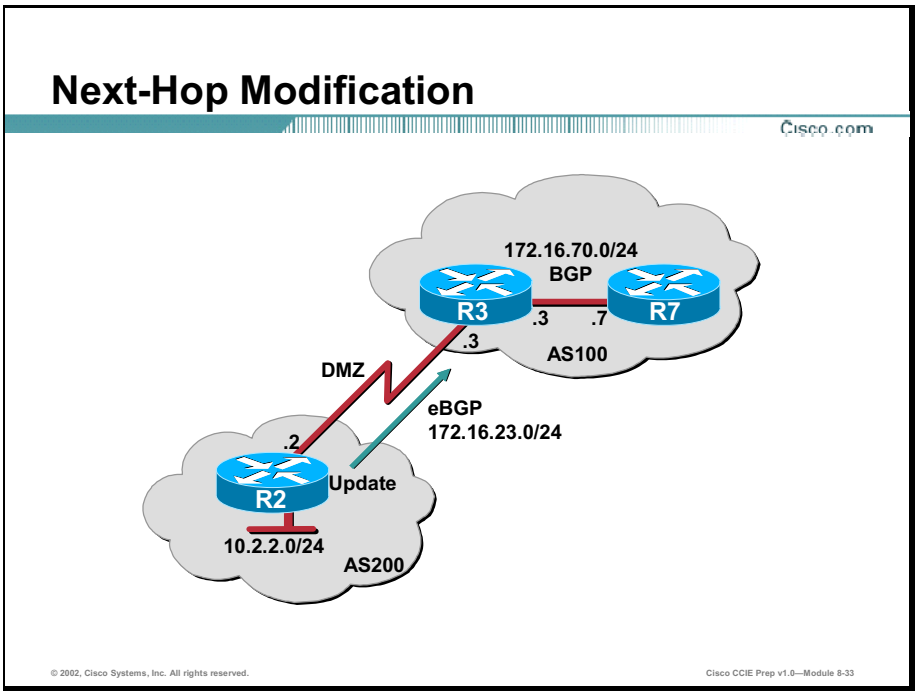

neighbor *{ip-address | peer-group-name}* **next-hop-self**

RFC1771 Section 5.1.3 NEXT_HOP

"When a BGP speaker advertises the route to a BGP speaker located in its own autonomous system, the advertising speaker shall not modify the NEXT_HOP attribute associated with the route."

Basically, this is saying if a border router receives an update from its eBGP neighbor, it will send an update to its iBGP peers with the next hop attribute unchanged. Now, look at an example.

In this example, neither AS is advertising the Demilitarized Zone (DMZ) network (172.16.23.0/24) in its IGP. R2 is configured as an eBGP peer with R3. R3 is configured as an iBGP peer with R7. R2 is advertising the network 10.2.2.0/24 to AS100.

```
R2(config)# router bgp 200
R2(config-router)# neighbor 172.16.23.3 remote-as 100
R2(config-router)# network 10.2.2.0 mask 255.255.255.0

R3(config)# router bgp 100
R3(config-router)# neighbor 172.16.23.2 remote-as 200
R3(config-router)# neighbor 172.16.70.7 remote-as 100

R7(config)# router bgp 100
R7(config-router)# neighbor 172.16.70.3 remote-as 100
```

Look at the BGP table on R3:

```
R3#show ip bgp
BGP table version is 2, local router ID is 172.16.70.3
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
*> 10.2.2.0/24      172.16.23.2              0             0 200 i
```

Here, you see the network 10.2.2.0 is reachable via AS200, using the next hop 172.16.23.2, which is what you would expect to see. Because you learned this route via eBGP and you satisfied the requirement that the next hop is known to the IGP, you can place the route to 10.2.2.0/24 in the routing table. You can verify this by the *> status codes on the 10.2.2.0/24 network. This can be verified with the **show** command of the routing table.

```
R3#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 not set
     172.16.0.0/24 is subnetted, 2 subnets
C       172.16.23.0 is directly connected, Serial1
C       172.16.70.0 is directly connected, Serial0
     10.0.0.0/24 is subnetted, 1 subnets
B       10.2.2.0 [20/0] via 172.16.23.2, 00:15:41
```

Now, look at R7's BGP table.

```
R7#show ip bgp
BGP table version is 1, local router ID is 172.16.70.7
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
* i10.2.2.0/24      172.16.23.2              0    100      0 200 i
```

Here on R7, you have received the route to 10.2.2.0/24 in our BGP table, and you see the same next hop information. Even though this route is in our BGP table, it has not been inserted in the IP routing table. Notice the ">" best status code is missing after the *. You can verify this by performing a **show ip route** on R7.


```
R7#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 not set
     172.16.0.0/24 is subnetted, 1 subnets
C       172.16.70.0 is directly connected, Serial0/1
```

Two important items must be met before this route 10.2.2.0/24 can be placed in the IP routing table. The synchronization rule must be met, and the next hop requirement must be met. Neither of which are met on R7 at this time.

The next hop requirement is not met because R7 does not have a route to 172.16.23.0/24. This will be fixed when you add the next-hop-self option to R3 for its neighbor R7.

The synchronization rule can be met by either redistributing BGP into the IGP, or by using the no synchronization option on R7.

Now, it is time to concentrate on the next hop requirement, which has not been satisfied. You have a few options that would allow you to pass the next hop requirement:

- Advertise the DMZ network in our IGP
- Create a static route to the DMZ network
- Use the next-hop-self attribute on R3

You will issue the **next-hop-self** command on R3 in its neighbor command to R7.

```
R3(config)# router bgp 100
R3(config-router)# neighbor 172.16.70.7 next-hop-self
```

Since you modified the BGP attributes to a neighbor, you must also clear the BGP connections to have the new policy updated. Issue the following command:

```
R3# clear ip bgp *
```

Now, look at R7's BGP table again.

```
R7#show ip bgp
BGP table version is 2, local router ID is 172.16.70.7
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
* i10.2.2.0/24      172.16.70.3              0    100      0 200 i
```

Notice that the next hop is R3's serial interface 172.16.70.3, just as you expected, but you are still not placing this route in the routing table. This is because the IGP is not aware of a route to 10.2.2.0/24. If you issue the command **no synchronization** on R7, this will by-pass the synchronization requirement and install the route in the routing table.

```
R7(config)# router bgp 100
R7(config-router)# no synchronization
```

Now that both the next hop requirement and the synchronization requirement have been met, R7 should install the route to 10.2.2.0/24 in its IP routing table. First, clear the BGP connection to allow the no synchronization policy to take effect. Then look at R7's BGP table again.

```
R7#clear ip bgp *
R7#show ip bgp
BGP table version is 2, local router ID is 172.16.70.7
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
*>i10.2.2.0/24      172.16.70.3              0    100      0 200 i
```

Now, you should have the route to 10.2.2.0/24 in R7's IP routing table.

```
R7#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 not set
     172.16.0.0/24 is subnetted, 1 subnets
C       172.16.70.0 is directly connected, Serial0/1
     10.0.0.0/24 is subnetted, 1 subnets
B       10.2.2.0 [200/0] via 172.16.70.3, 00:11:39
```

Normally, unless you are advertising your DMZ into the IGP, you will be setting the next-hop attribute on your border router going to each of your iBGP peers.

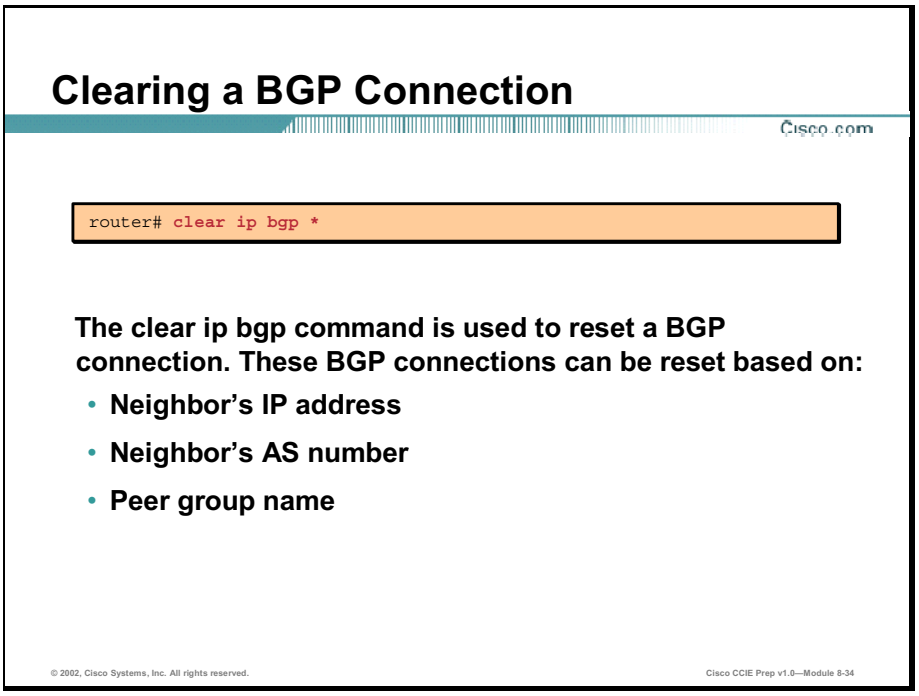


```
clear ip bgp {*|address} [soft [in|out]]
```

The clear ip bgp command is used to reset a BGP connection. BGP connections can be reset based on:

- Neighbor's IP address
- Neighbor's AS number
- Peer group name

You must reset your BGP connections when any of the following have been modified or added to:

- BGP route map
- BGP distribute list
- BGP weight
- BGP administrative distance
- BGP timers
- BGP access list

To clear or reset all BGP connections on a router, issue the command **clear ip bgp ***.

The problem with this command is it transitions the BGP peer from established to idle, then rebuilds the relationship along with any new policies. In a production environment, this can cause a significant delay and a long down time for the clients.

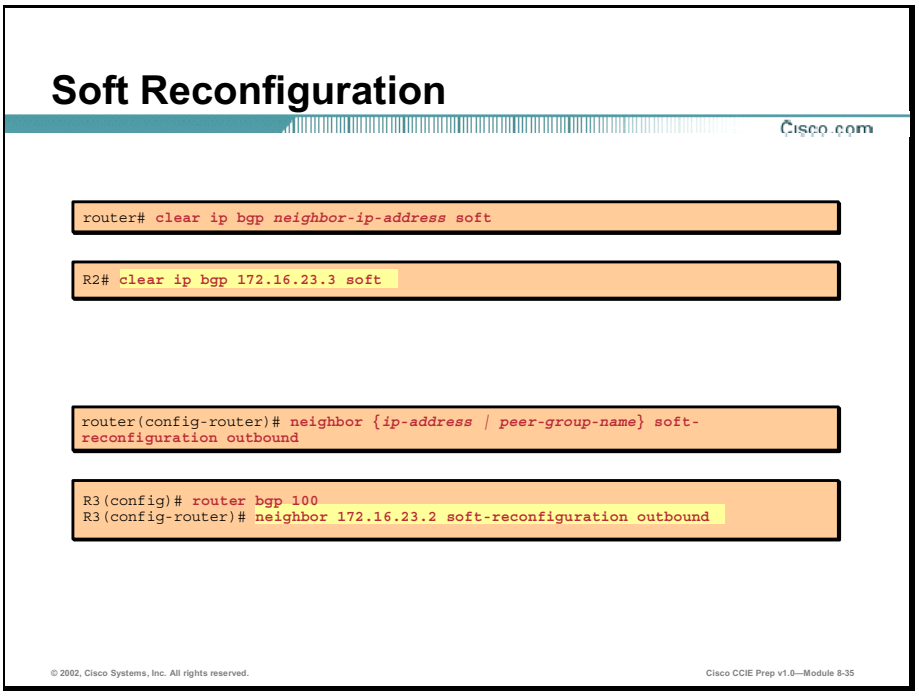

clear ip bgp *neighbor-ip-address* **soft**

You can eliminate resetting the BGP connections if you perform a soft reconfiguration. With a soft reconfiguration, you do not reset the connection. Instead, you resend all the routing updates.

`R2# clear ip bgp 172.16.23.3 soft`

When you issue this command, you must also configure the BGP peer (172.16.23.3 = R3) to allow this soft reconfiguration. To allow R2 to perform a soft reconfiguration on R3, you would issue the following command on R3.

neighbor {*ip-address* | *peer-group-name*} **soft-reconfiguration outbound**

```
R3(config)# router bgp 100
R3(config-router)# neighbor 172.16.23.2 soft-reconfiguration outbound
```

Summary

This section summarizes the key points discussed in this lesson.

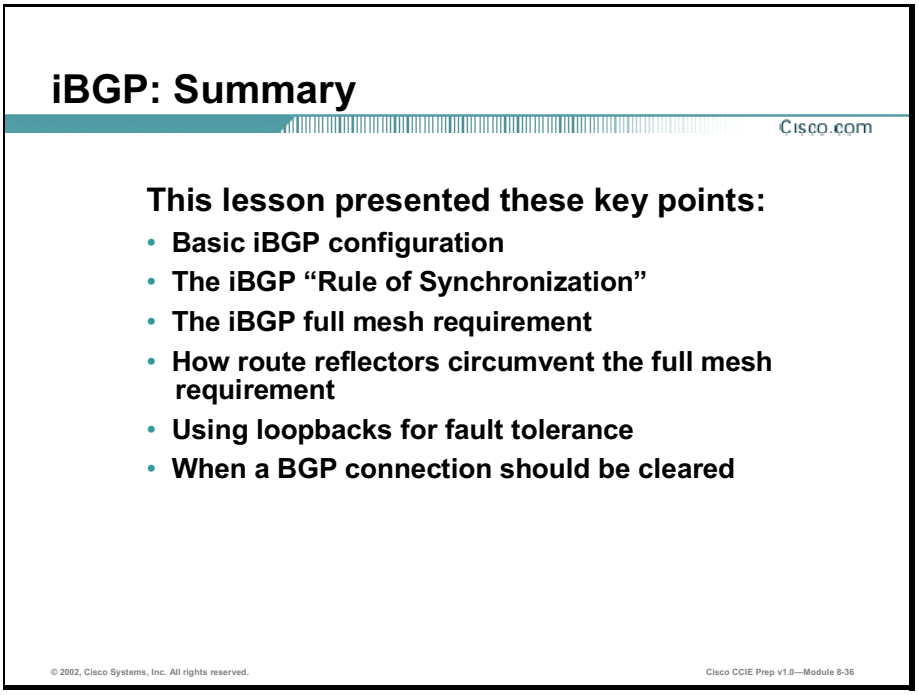

Next Steps

After completing this lesson, go to:

- eBGP Configuration

References

For additional information, refer to these resources:

- *Internet Routing Architectures* by Sam Halabi

Lesson Assessment (Quiz)

Q1) When your BGP autonomous system ID matches that of your BGP neighbor, what is this considered to be?

A) An EGP relationship

B) External BGP

C) Internal BGP

D) An IGP relationship

Q2) When running a full mesh iBGP with 10 BGP speakers, how many total peer connections are required?

A) One

B) Four

C) Forty Five

D) Ninety

Q3) Using laymen's terms, what does the iBGP synchronization rule state?

A) Any and all routes must be synchronized with the IGP before being placed in the BGP table.

B) Any and all routes must be synchronized with the EGP before being placed in the IP routing table.

C) All BGP peers must have the same (synchronized) BGP table before routes can be placed in the IP routing table.

D) Do not advertise a route if the IGP does not have it in its routing table.

Q4) When creating route reflection for a specific client, on which iBGP peer should the command(s) be placed?

A) The server

B) The client

C) All iBGP peers

D) The hub router in the iBGP

Q5) When you have modified an access list used with your BGP neighbor statement, which action would be performed next?

A) Clear the route map

B) Clear the iBGP connections

C) Reload the router

D) Apply the access list to an interface

eBGP Configuration

Overview

The Border Gateway Protocol (BGP) is a distance-vector routing protocol that is used to exchange routing information between different Autonomous Systems (AS). This lesson will describe how to configure basic External BGP (eBGP) as well as more advanced options such as eBGP multihop, confederations, and communities.

Importance

BGP version 4 is the current exterior routing protocol used on the Internet. When connecting a multi-homed enterprise to the Internet, it is vital that the administrator know BGP properties, configurations and troubleshooting tips.

Objectives

Upon completing this lesson, you will be able to:

- Configure basic eBGP peer relationships
- Configure advanced eBGP options such as multihop
- Describe BGP confederations and how to configure them
- Describe BGP communities and how to configure them

Learner Skills and Knowledge

To fully benefit from this lesson, you must have these prerequisite skills and knowledge:

- Passed the Cisco Certified Internetworking Expert (CCIE) written qualification exam
- Knowledge of Transfer Control Protocol /Internet Protocol (TCP/IP), access control lists, and Interior Gateway Protocols (IGPs)

Outline

This lesson includes these sections:

- Overview
- eBGP Basic Configuration
- eBGP Advanced Configuration
- Advanced Configuration Options
- Communities
- Summary
- Lesson Assessment (Quiz)

eBGP Basic Configuration

This section discusses basic eBGP and how it is configured.

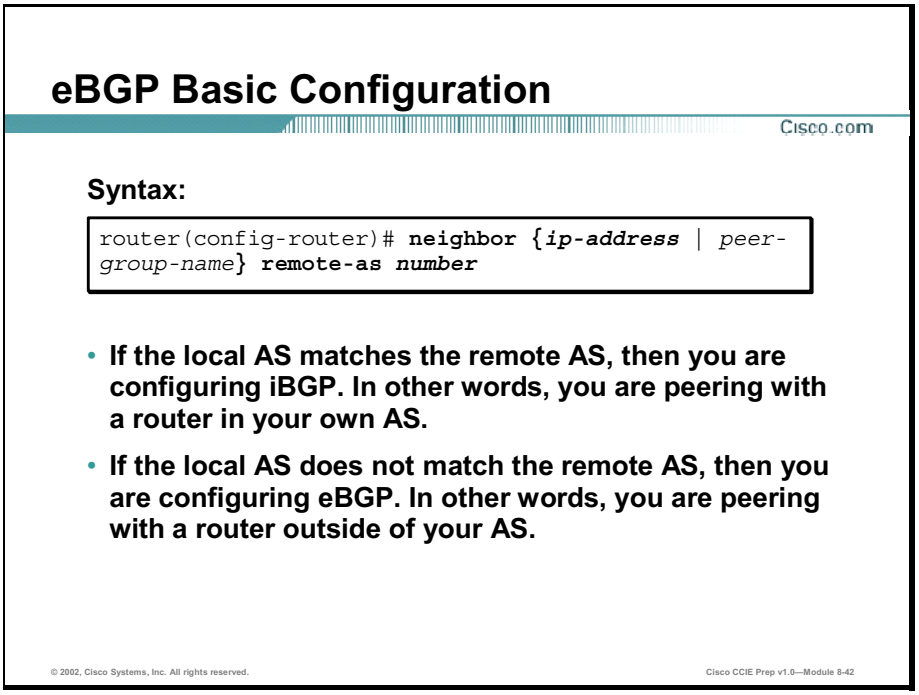

There are two primary commands required to configure an External BGP (eBGP) neighbor relationship.

router bgp *<AS-number>*

neighbor *{ip-address | peer-group-name}* **remote-as** *AS-number*

As you can see, the commands required to configure an eBGP neighbor relationship are identical to the commands required to configure an Internal BGP (iBGP) neighbor relationship. The distinguishing difference between an iBGP or eBGP neighbor relationship is between the local Autonomous System (AS) and the remote AS.

If the local AS matches the remote AS, then you are configuring iBGP. You are peering with a router in your own AS.

If the local AS does not match the remote AS, then you are configuring eBGP. In other words, you are peering with a router outside of your own AS.

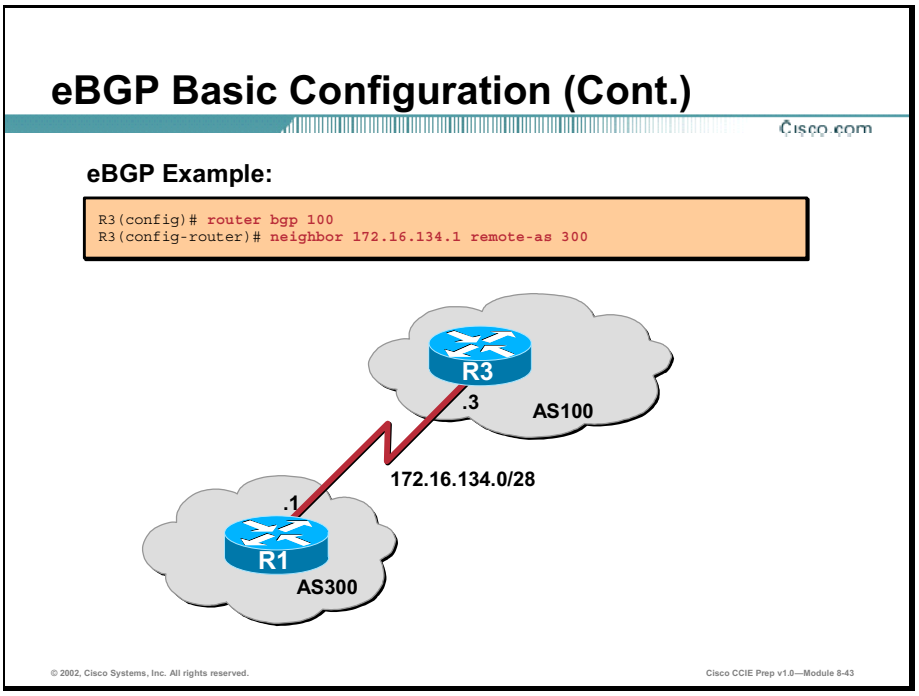

When you use iBGP in the AS, you could be acting as a transit area for two or more different autonomous systems. When you use eBGP, you are connecting to another AS to obtain routes to their resources, or resources they have learned about.

eBGP Advanced Configuration

This section will describe when to use eBGP multihop and how to configure it.

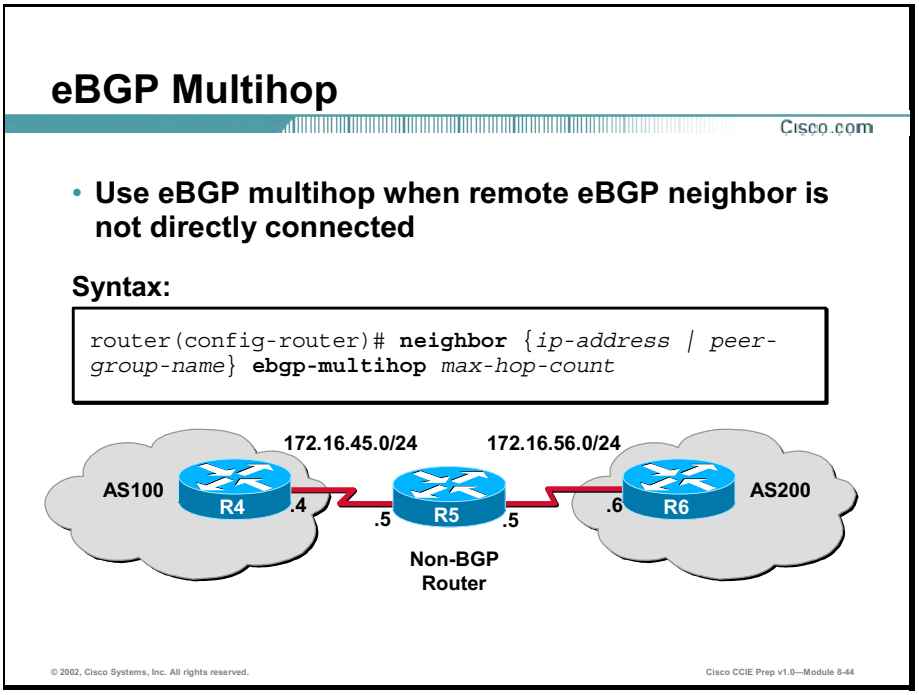

Usually, the two eBGP speakers are directly connected (for example, over a Wide Area Network (WAN) connection). Sometimes, they cannot be directly connected, such as the case when a router that does not use BGP is in between the two neighbors that wish to form an eBGP neighbor relationship. In this special case, the **neighbor ebgp-multihop** router configuration command is used. Without this command, an eBGP neighbor relationship with a non-directly connected neighbor will never form. Remember, this command is only used with eBGP, not iBGP.

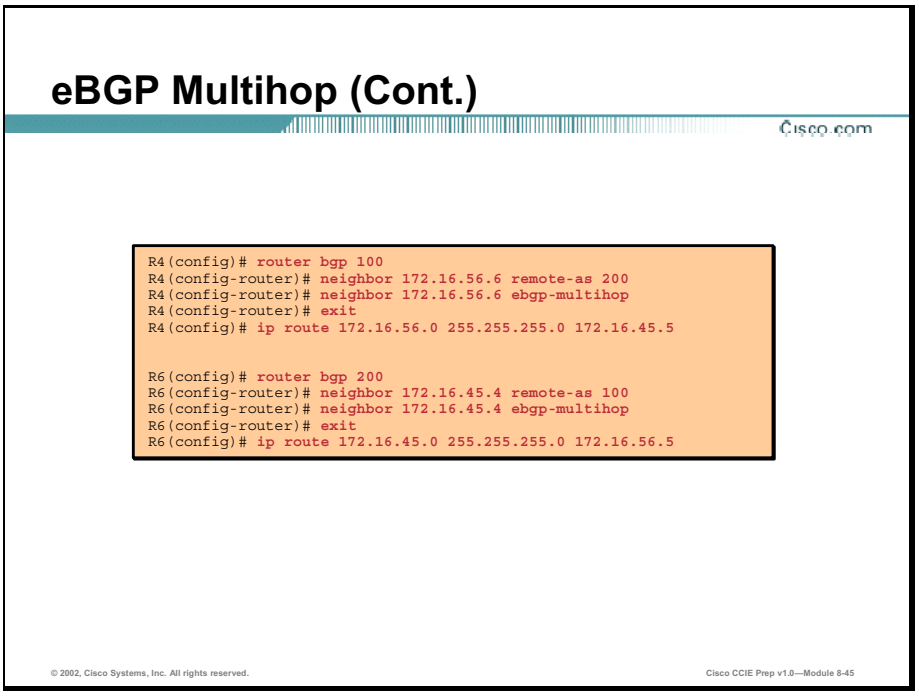

neighbor *{ip-address | peer-group-name}* **ebgp-multihop** *max-hop-count*

Look at the previous example where R5 is a non-BGP speaking router. R4 and R6 wish to form an eBGP neighbor relationship.

Notice that both routers reference their external neighbor by an Internet Protocol (IP) address that is not directly connected. A requirement of BGP is reachability, so an **ebgp-multihop** configuration must include static routes or must enable an IGP so that the neighbors can reach each other.

eBGP Load Balancing

Cisco.com

- To load balance use eBGP-multihop in conjunction with update-source loopback

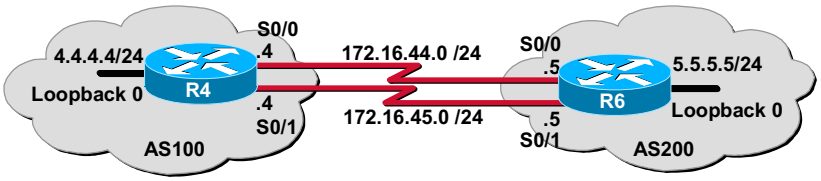

It is possible to have a situation where an **ebgp-multihop** router configuration is useful together with loopbacks. Using these two features enables you to perform load balancing between two autonomous systems over parallel links.

```
R4(config)# interface loopback 0
R4(config-if)# ip address 4.4.4.4 255.255.255.0
R4(config-if)# router bgp 100
R4(config-router)# neighbor 5.5.5.5 remote-as 200
R4(config-router)# neighbor 5.5.5.5 ebgp-multihop
R4(config-router)# neighbor 5.5.5.5 update-source loopback 0
R4(config-router)# network 4.4.4.0 mask 255.255.255.0
R4(config-router)# exit
R4(config)# ip route 5.5.5.0 255.255.255.0 172.16.44.5
R4(config)# ip route 5.5.5.0 255.255.255.0 172.16.45.5

R5(config)# interface loopback 0
R5(config-if)# ip address 5.5.5.5 255.255.255.0
R5(config-if)# router bgp 200
R5(config-router)# neighbor 4.4.4.4 remote-as 100
R5(config-router)# neighbor 4.4.4.4 ebgp-multihop
R5(config-router)# neighbor 4.4.4.4 update-source loopback 0
R5(config-router)# network 5.5.5.0 mask 255.255.255.0
R5(config-router)# exit
R5(config)# ip route 4.4.4.0 255.255.255.0 172.16.44.4
```

```
R5(config)# ip route 4.4.4.0 255.255.255.0 172.16.45.4
```

Look at R4's BGP route table:

```
R4# show ip bgp
BGP table version is 3, local router ID is 172.16.45.4
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
*> 4.4.4.0/24       0.0.0.0                  0         32768 i
*> 5.5.5.0/24       5.5.5.5                  0      0    200 i
```

Notice that the route to 5.5.5.0/24 is via 5.5.5.5 as the next hop. Look at what the IP routing table looks like for this route:

```
R4# 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 not set
     4.0.0.0/24 is subnetted, 1 subnets
C       4.4.4.0 is directly connected, Loopback0
     5.0.0.0/24 is subnetted, 1 subnets
S       5.5.5.0 [1/0] via 172.16.44.5
                [1/0] via 172.16.45.5
     172.16.0.0/24 is subnetted, 2 subnets
C       172.16.44.0 is directly connected, Serial1/0
C       172.16.45.0 is directly connected, Serial1/1
```

The **neighbor ebgp-multihop** and **neighbor update-source** router configuration commands have the effect of making the loopback interface the next hop for eBGP, which allows load balancing to occur. You can use static routes to introduce two equal-cost paths to the destination. (The same effect could also be accomplished by using an IGP.) R4 can reach the next hop of 5.5.5.5 in two ways: via 172.16.44.5 and via 172.16.45.5. Also, R5 can reach the next hop of 4.4.4.4 in two ways: via 172.16.44.4 and via 172.16.45.4.

Advanced Configuration Options

This section will discuss BGP confederations and how to configure them.

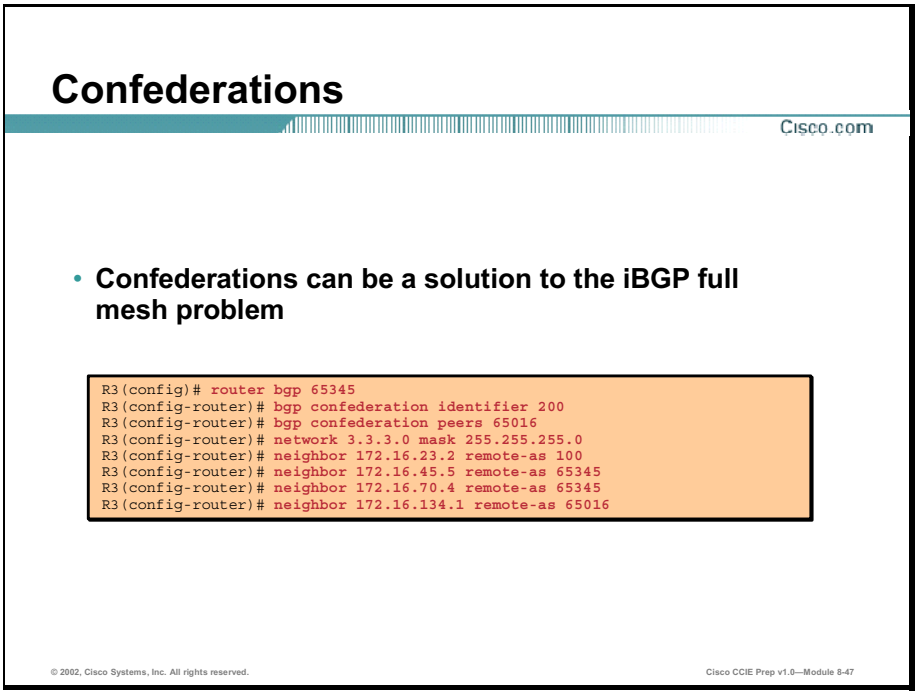

One way to reduce the iBGP mesh is to divide an autonomous system into multiple sub-autonomous systems and group them into a single confederation. To the outside world, the confederation looks like a single autonomous system. Each autonomous system is fully meshed within, and has a few connections to other autonomous systems in the same confederation. Even though the peers in different autonomous systems have eBGP sessions, they exchange routing information as if they were iBGP peers. Specifically, the next hop, Multi-Exit Discriminator (MED), and local preference information is preserved. This feature allows you to retain a single IGP for all of the autonomous systems.

To configure a BGP confederation, you must specify a confederation identifier. To the outside world, the group of autonomous systems will look like a single autonomous system with the confederation identifier as the autonomous system number. To configure a BGP confederation identifier, use the following commands in router configuration mode:

bgp confederation identifier *AS-number*

bgp confederation peers *AS-numbers*

The **bgp confederation identifier** configures a BGP confederation.

The **bgp confederation peers** specifies the autonomous systems that belong to the confederation.

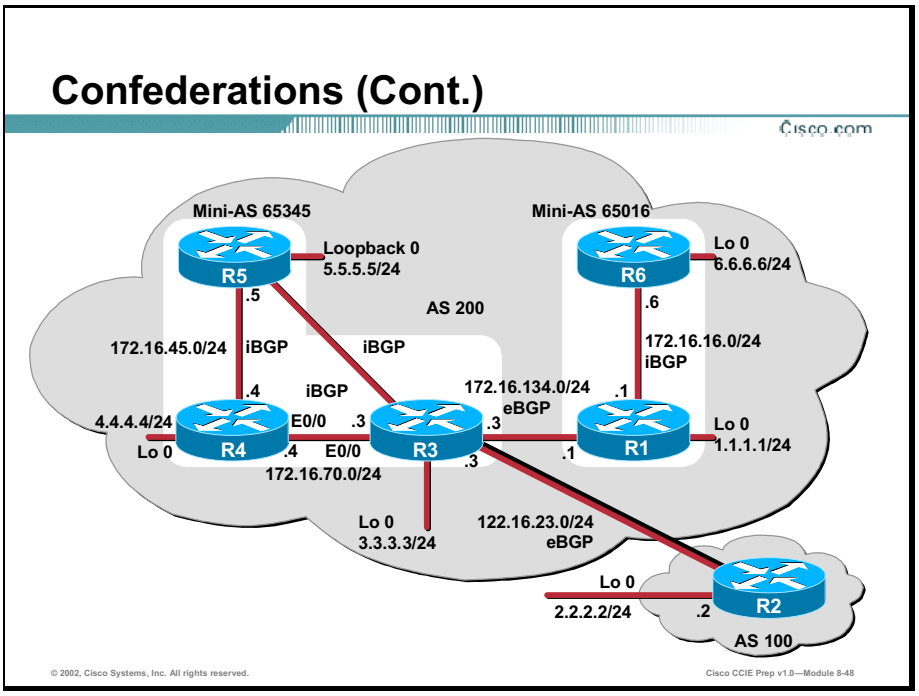

In this scenario, AS200 is running a common IGP. Each router is advertising its loopback interface. AS200 is sub-divided into to mini-AS's 65345 and 65016.

When customers are multihomed to a single Internet Service Provider (ISP), the ISPs assign private Autonomous System (AS) numbers in order to conserve AS numbers. These private AS numbers come from the range 64512 to 65535.

A confederation is a technique for reducing the iBGP mesh inside the AS. In the diagram, AS 200 consists of multiple BGP speakers (although there might be other routers that are not configured for BGP). Without confederations, BGP would require the routers in AS 200 to be fully meshed. That is, each router would need to run iBGP with each of the other routers. Specifically:

- You use confederations to divide the AS into multiple sub-AS's and assign the sub-AS's to a confederation
- Each sub-AS is fully meshed, and iBGP is run among its members
- Each sub-AS has a connection to the other sub-AS's within the confederation

Remember, even though the sub-AS's have eBGP peers to AS's within the confederation, they exchange routing updates as if they were using iBGP.

Now look at the configuration for R3:

```
R3(config)# router bgp 65345
R3(config-router)# bgp confederation identifier 200
R3(config-router)# bgp confederation peers 65016
R3(config-router)# network 3.3.3.0 mask 255.255.255.0
```



```
R3(config-router)# neighbor 172.16.23.2 remote-as 100
R3(config-router)# neighbor 172.16.45.5 remote-as 65345
R3(config-router)# neighbor 172.16.70.4 remote-as 65345
R3(config-router)# neighbor 172.16.134.1 remote-as 65016
```

Analyze this configuration to see exactly what is happening.

```
R3(config)# router bgp 65345
```

This command is stating that R3 is part of AS 65345.

```
R3(config-router)# bgp confederation identifier 200
R3(config-router)# bgp confederation peers 65016
```

These commands work in conjunction with the **router bgp 65345** command. They are stating that this router is actually part of overall AS 200, but the sub-AS is 65345. That means you will perform iBGP peering with other routers in sub-AS 65345, not AS 200.

Other sub-AS's of AS 200 are defined with the **bgp confederation peers** command. In this case, only sub-AS 65016 is defined. If there were other sub-AS's they would also need to be defined here.

```
R3(config-router)# network 3.3.3.0 mask 255.255.255.0
```

This command states R3 will be advertising network 3.3.3.0 to its BGP peers.

```
R3(config-router)# neighbor 172.16.23.2 remote-as 100
R3(config-router)# neighbor 172.16.45.5 remote-as 65345
R3(config-router)# neighbor 172.16.70.4 remote-as 65345
R3(config-router)# neighbor 172.16.134.1 remote-as 65016
```

These commands define your neighbors.

Neighbor 172.16.23.2 is part of AS 100 an eBGP neighbor. Normal BGP processes occur with this neighbor.

Neighbors 172.16.45.5 and 172.16.70.4 are part of the mini-AS and standard iBGP processing occurs here.

Neighbor 172.16.134.1 is part of mini-AS 65016 R3's confederation peer. Even though you have an eBGP neighbor relationship and it is part of AS 200's confederation, normal iBGP processing will occur.

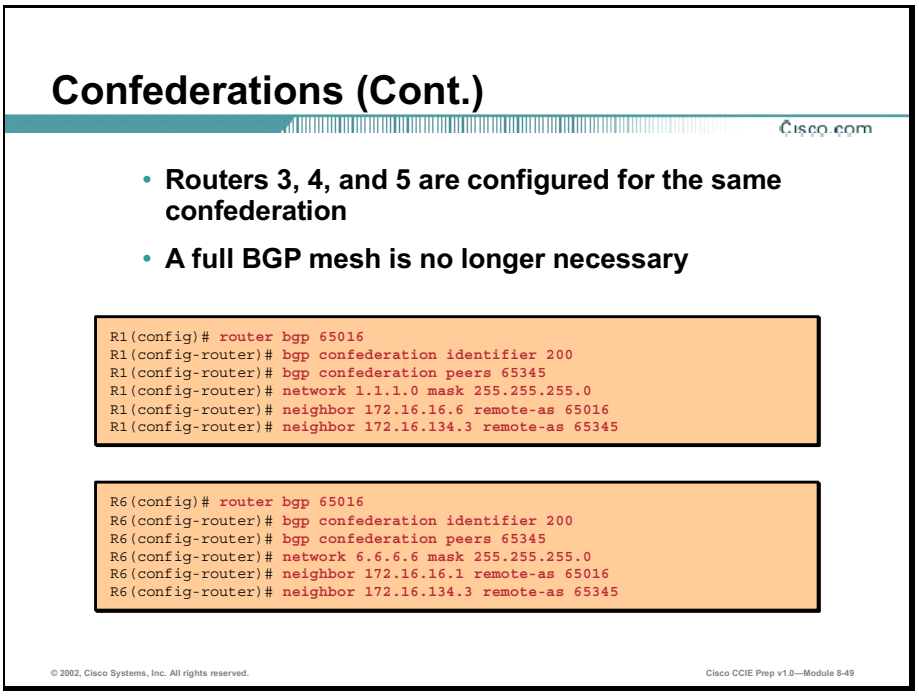

You can examine these two other configurations, specifically R6 and R1.

Notice that since R6 does not peer with a router in sub-AS 65345 the **bgp confederation peers** command is not required. Placing it in the configuration will not affect BGP. This is also true for R4 and R5 in AS 65345.

Since R1 peers with sub-AS 65345, the confederation peer statement is required here. Finally, view what the BGP table looks like on R1.

```
R1#show ip bgp
BGP table version is 13, local router ID is 1.1.1.1
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
*> 1.1.1.0/24       0.0.0.0                  0         32768 i
*> 2.2.2.0/24       172.16.23.2              0    100      0  (65345)100 i
*> 3.3.3.0/24       172.16.134.3             0    100      0  (65345) i
*> 4.4.4.0/24       172.16.70.4              0    100      0  (65345) i
*> 5.5.5.0/24       172.16.45.5              0    100      0  (65345) i
*>i6.6.6.0/24       172.16.16.6              0    100      0  I
```

Communities

This section will describe BGP communities and how to configure them.

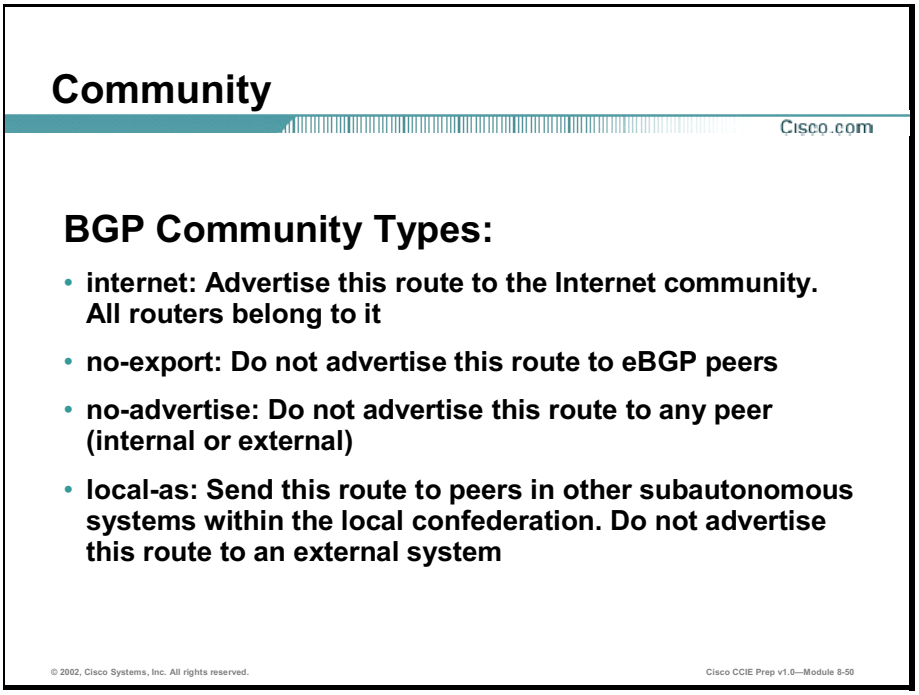

The *communities* attribute is a way to group destinations into communities and apply routing decisions based on the communities. This method simplifies the configuration of a BGP speaker that controls distribution of routing information.

A community is a group of destinations that share some common attribute. Each destination can belong to multiple communities. Autonomous system administrators can define to which communities a destination belongs. By default, all destinations belong to the general Internet community. The community is carried as the communities attribute.

The *communities* attribute is an optional, transitive, global attribute in the numerical range from 1 to 4,294,967,200. Along with Internet community, there are a few predefined, well-known communities, as follows:

internet—Advertise this route to the Internet community. All routers belong to it.
no-export—Do not advertise this route to eBGP peers.
no-advertise—Do not advertise this route to any peer (internal or external).
local-as—Send this route to peers in other subautonomous systems within the local confederation. Do not advertise this route to an external system.

Based on the community, you can control which routing information to accept, prefer, or distribute to other neighbors. A BGP speaker can set, append, or modify the community of a route when you learn, advertise, or redistribute routes. When routes are aggregated, the resulting aggregate has a *communities* attribute that contains all communities from all the initial routes.

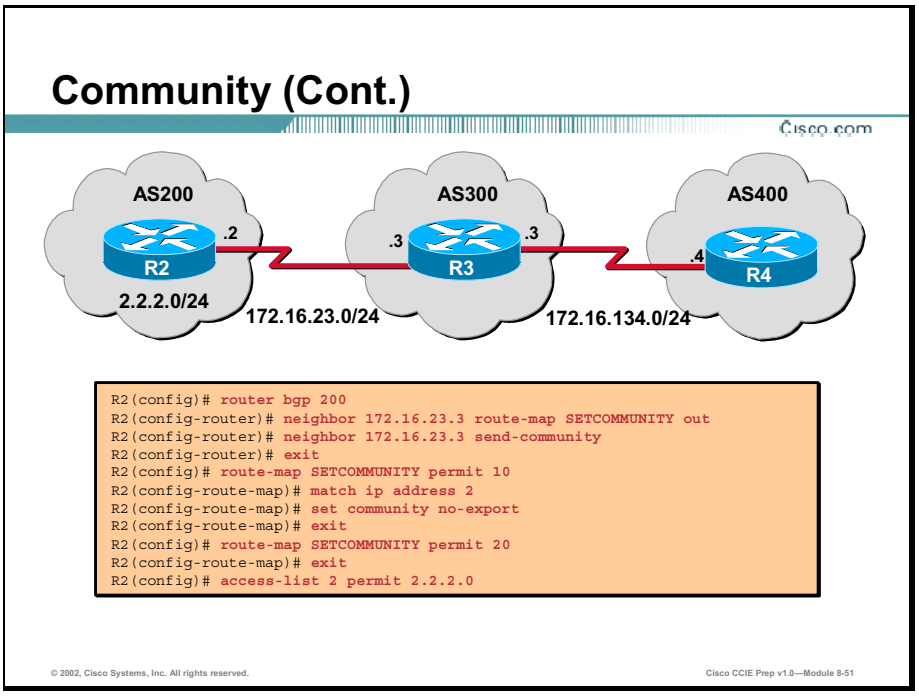

Now look at a scenario in which AS 200 needs to share information on network 2.2.2.0/24 to AS 300, but not any other AS outside of AS 300.

In this scenario, R2 advertises networks to its eBGP neighbor R3 via a route-map. The route-map stipulates that network 2.2.2.0/24 should be advertised with the no-export community attribute set. All other networks will not set a community attribute.

Here is R2's configuration to accomplish this:

```
R2(config)# router bgp 200
R2(config-router)# neighbor 172.16.23.3 route-map SETCOMMUNITY out
R2(config-router)# neighbor 172.16.23.3 send-community
R2(config-router)# exit
R2(config)# route-map SETCOMMUNITY permit 10
R2(config-route-map)# match ip address 2
R2(config-route-map)# set community no-export
R2(config-route-map)# exit
R2(config)# route-map SETCOMMUNITY permit 20
R2(config-route-map)# exit
R2(config)# access-list 2 permit 2.2.2.0
```

Analyze this configuration to see exactly what is happening.

```
R2(config-router)# neighbor 172.16.23.3 route-map SETCOMMUNITY out
```


This command will run the route map SETCOMMUNITY before advertising routes outbound to R3.

```
R2(config-router)# neighbor 172.16.23.3 send-community
```

By default, BGP does not send the community attribute; this command is required in order to do so. If you omit this command, even if you set a community attribute, it will be stripped before sending routes to R3.

```
R2(config)# route-map SETCOMMUNITY permit 10
R2(config-route-map)# match ip address 2
R2(config-route-map)# set community no-export
R2(config-route-map)# exit
R2(config)# route-map SETCOMMUNITY permit 20
```

The route map is used to modify the policies or in this case the community attribute of certain routes. In the **permit 10** section, you are matching any address in access list 2, which in this case is the 2.2.2.0/24 network. If the address is matched, then you set the community attribute to no-export and advertise the route. The **permit 20** command is very important. Remember at the end of each route map if no condition for a route is met, that route will not be advertised. The **permit 20** statement alone indicates "permit all other routes to be advertised."

Now, view the no-export community on R3 to make sure you have received it for the 2.2.2.0/24 network.

```
R3#show ip bgp community no-export
BGP table version is 3, local router ID is 3.3.3.3
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
*> 2.2.2.0/24       172.16.23.2           0         0   200    i
```

Summary

This section summarizes the key points discussed in this lesson.

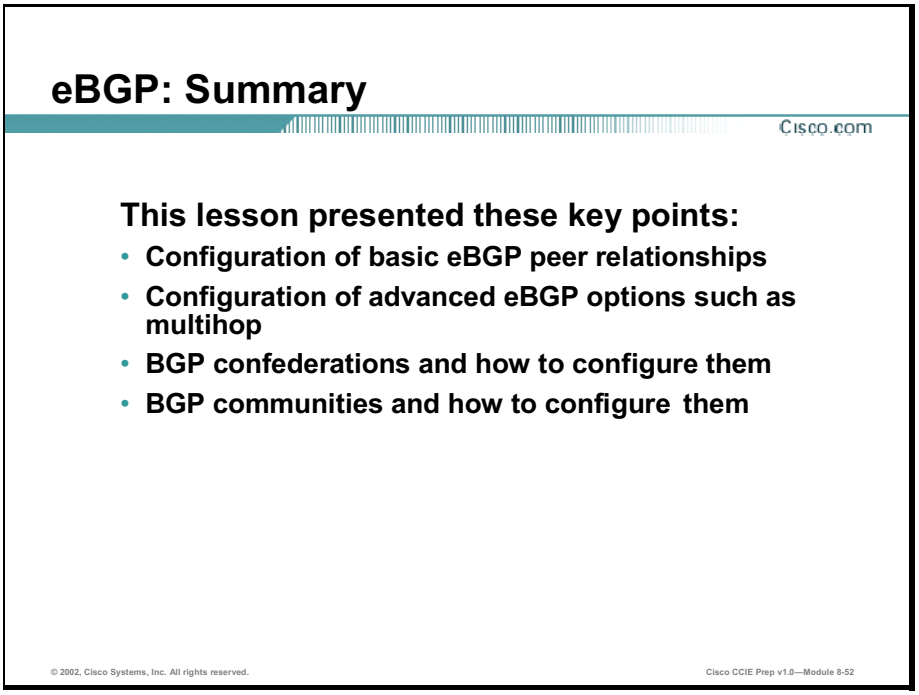

Next Steps

After completing this lesson, go to:

- Advertising Networks

References

For additional information, refer to these resources:

- *Internet Routing Architectures* by Sam Halabi

Lesson Assessment (Quiz)

Q1) True or False. In most situations iBGP neighbors are not directly connected while eBGP neighbors are.

A) True

B) False

Q2) Which of the following lessens the full mesh requirement?

A) eBGP multihop

B) confederations

C) communities

D) using loopback interfaces

E) route reflectors

Q3) Which of the following is used to simplify the configuration of a BGP speaker that controls distribution of routing information?

A) eBGP multihop

B) confederations

C) communities

D) using loopback interfaces

Q4) Which of the following communities is set by default on all destinations?

A) internet

B) no-export

C) no-advertise

D) local-as

Q5) After modifying the community being sent to a neighbor, which of the following commands must also be issued?

A) neighbor <ip-address> send-community

B) neighbor <ip-address> advertise-community

C) clear ip bgp

D) neighbor <ip-address> receive-community

Advertising Networks

Overview

The Border Gateway Protocol (BGP) is a distance-vector routing protocol that is used to exchange routing information among different Autonomous Systems (AS). This lesson will discuss when and how to advertise networks via BGP.

Importance

BGP version 4 is the current exterior routing protocol used on the Internet. When connecting a multi-homed enterprise to the Internet, it is vital that the administrator know BGP properties, configurations and troubleshooting tips.

Objectives

Upon completing this lesson, you will be able to:

- Describe the advertising methods available when using BGP
- Configure static route redistribution
- Configure dynamic route redistribution
- Configure route advertisement using the **network** command

Learner Skills and Knowledge

To fully benefit from this lesson, you must have these prerequisite skills and knowledge:

- Passed the Cisco Certified Internetworking Expert (CCIE) written qualification exam
- Knowledge of Transfer Control Protocol/Internet Protocol (TCP/IP), access control lists, and Interior Gateway Protocols (IGPs)

Outline

This lesson includes these sections:

- Overview
- Advertising Methods
- Redistributing Static Routes
- Redistributing Dynamic Routes
- Using the Network Command
- Summary
- Lesson Assessment (Quiz)

Advertising Methods

This section will discuss methods to use for route advertisement.

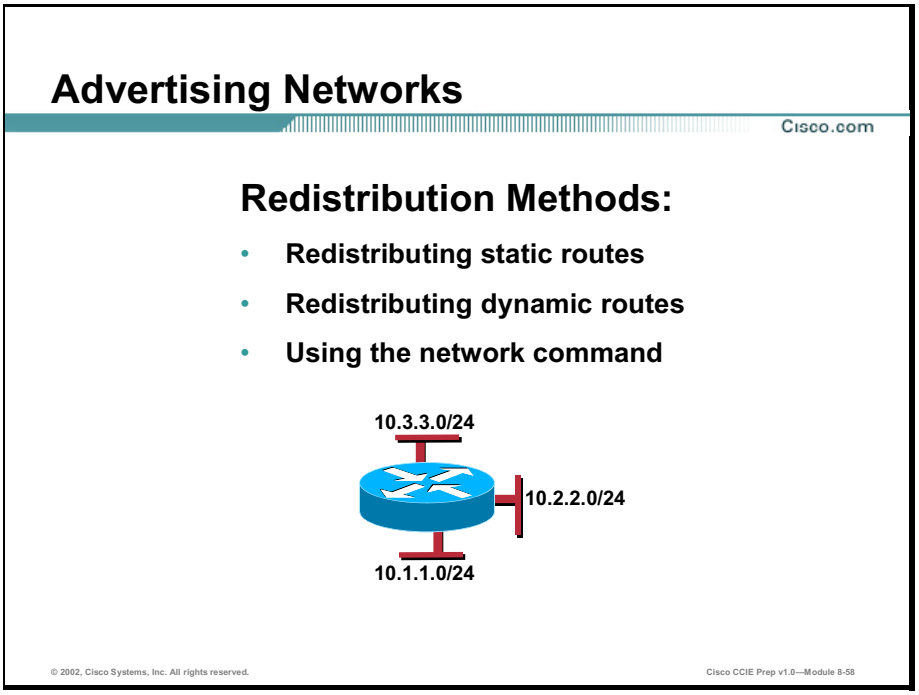

A network that resides within an Autonomous System (AS) is said to originate from that network. To inform other ASs about its networks, the AS advertises them. Border Gateway Protocol (BGP) provides three ways for an AS to advertise the networks that it originates:

- Redistributing static routes
- Redistributing dynamic routes
- Using the network command

It is important to remember that routes advertised by the techniques described in this section are advertised *in addition* to other BGP routes that a BGP-configured router learns from its internal and external neighbors. BGP always passes on information that it learns from one peer to other peers. The difference is that routes generated by the **network** and **redistribute** router configuration commands specify the AS of the router as the originating AS for the network.

Redistributing Static Routes

This section will discuss how to perform static route redistribution.

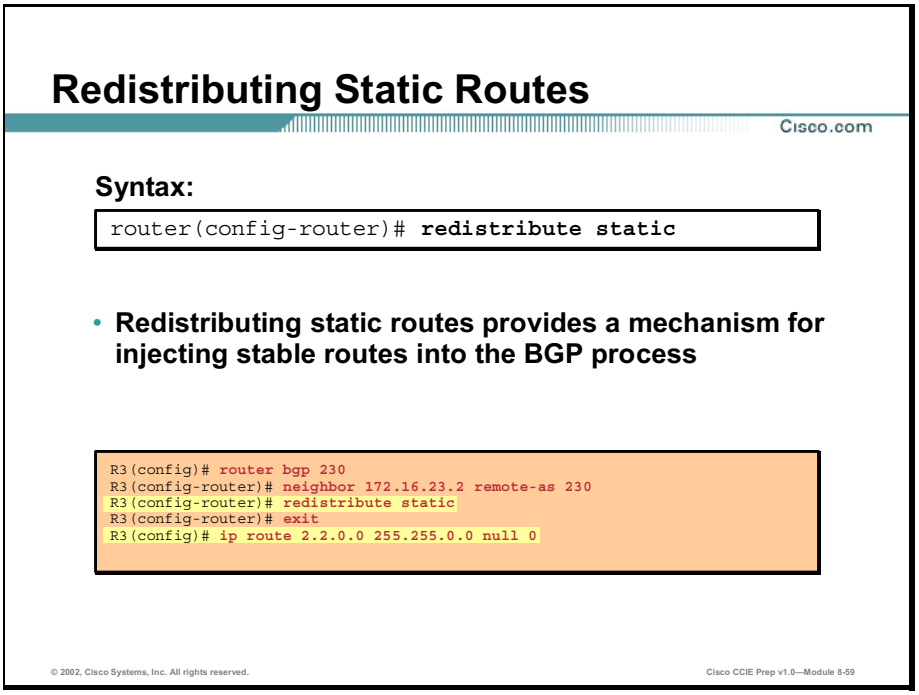

ip route *<destination-network> <mask> {next-hop | interface}*

One way to advertise that a network or a subnet originates from an AS is to redistribute static routes into BGP. The only difference between advertising a static route and advertising a dynamic route is that when you redistribute a static route, BGP sets the origin attribute of updates for the route to Incomplete.

To configure R3 to originate network 2.2.0.0/16 into BGP, use these commands:

```
R3(config)# router bgp 230
R3(config-router)# neighbor 172.16.23.2 remote-as 230
R3(config-router)# redistribute static
R3(config-router)# exit
R3(config)# ip route 2.2.0.0 255.255.0.0 null 0
```

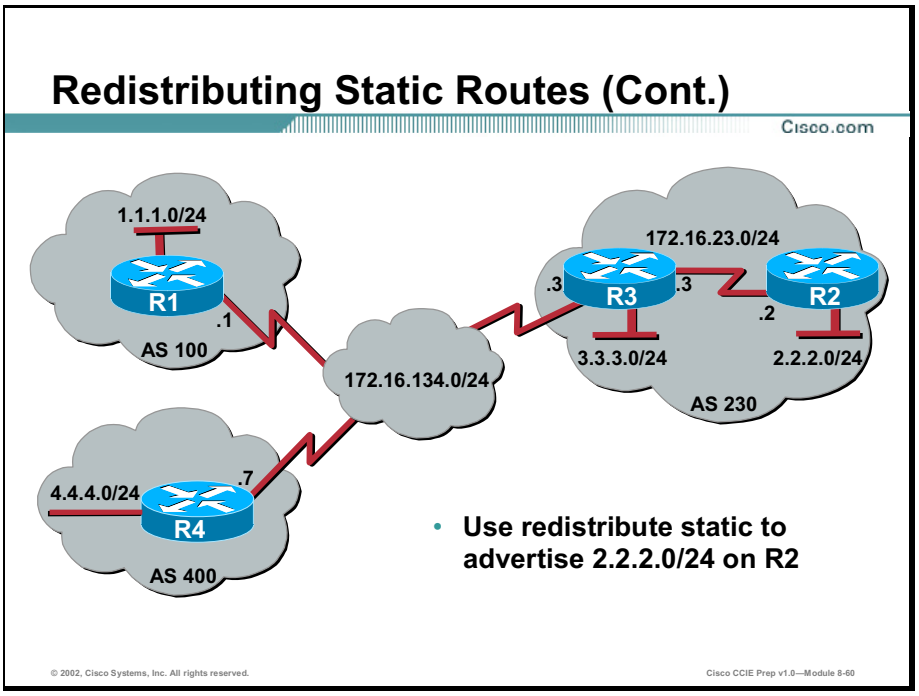

The **redistribute** router configuration command and the **static** keyword cause all static routes to be redistributed into BGP.

The **ip route** global configuration command establishes a static route for network 2.2.0.0. In theory, the specification of the null 0 interface would cause a packet destined for network 2.2.0.0 to be discarded. In practice, there will be a more specific match for the packet than 2.2.0.0, and the router will send it out the appropriate interface. Redistributing a static route is the best way to advertise a supernet because it prevents the route from flapping.

Redistributing Dynamic Routes

This section will discuss how to perform dynamic route redistribution.

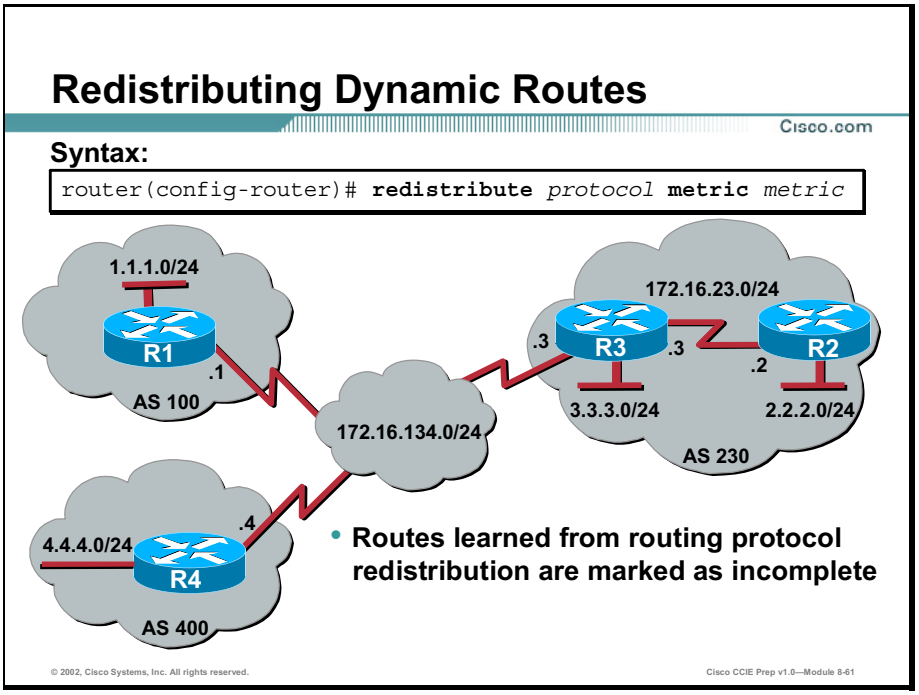

redistribute *protocol* **metric** *metric*

Another way to advertise networks is to redistribute dynamic routes. Typically, you redistribute Interior Gateway Protocol (IGP) routes (such as Enhanced Interior Gateway Routing Protocol (EIGRP), IGRP, Intermediate System to Intermediate System (IS-IS), Open Shortest Route First (OSPF), and Routing Information Protocol (RIP) routes) into BGP. Some of your IGP routes might have been learned from BGP, so you need to use access lists to prevent the redistribution of routes back into BGP.

Routers R2 and R3 are running iBGP. R3 is learning 4.4.4.0/24 via BGP, and redistributing 4.4.4.0/24 back into Enhanced IGRP. The following commands configure R3:

```
R3(config)# router eigrp 10
R3(config-router)# network 3.3.3.0
R3(config-router)# redistribute bgp 230
R3(config-router)# redistributed connected
R3(config-router)# default-metric 1000 100 250 1 1500
R3(config-router)# exit
R3(config)# router bgp 230
R3(config-router)# neighbor 172.16.134.1 remote-as 100
R3(config-router)# neighbor 172.16.23.2 remote-as 230
R3(config-router)# neighbor 172.16.134.1 distribute-list 1 out
R3(config-router)# redistribute eigrp 10
```

```
R3(config-router)# exit
R3(config)# access-list 1 permit 2.2.2.0 0.0.0.255
```

The **redistribute** router configuration command with the **eigrp** keyword redistributes Enhanced IGRP routes for process ID 10 into BGP.

Note	Normally, distributing BGP into IGP should be avoided because too many routes would be injected into the AS.

The **neighbor distribute-list** router configuration command applies access list 1 to outgoing advertisements to the neighbor whose Internet Protocol (IP) address is 172.16.134.1 (that is, R1). Access list 1 specifies that network 2.2.2.0 is to be advertised. All other networks, such as network 4.4.4.0, are implicitly prevented from being advertised. The access list prevents network 4.4.4.0 from being injected back into BGP as if it originated from AS 230, and allows BGP to advertise network 2.2.2.0 as originating from AS 230.

Redistribution of dynamic routes requires careful use of access lists to prevent updates from being injected back into BGP. If possible, you should use the **network** command or redistribute static routes instead of redistributing dynamic routes.

Using the Network Command

This section will discuss how to advertise routes using the network command.

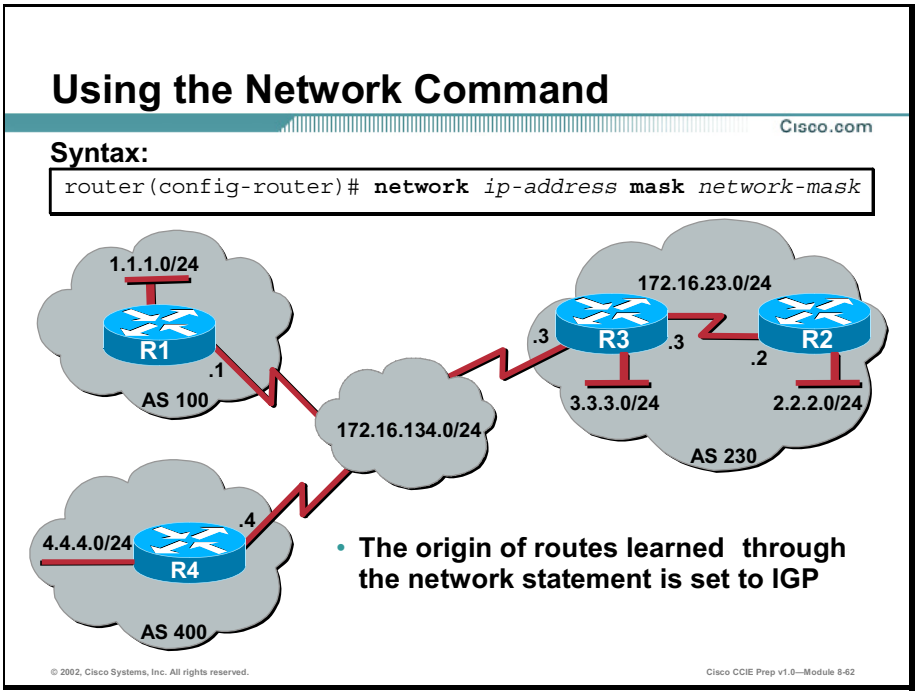

network *ip-address* **mask** *network-mask*

Another way to advertise networks is to use the **network** router configuration command. When used with BGP, the **network** command specifies the networks that the AS originates. By way of contrast, when used with an IGP such as Routing Information Protocol (RIP), the **network** command identifies the interfaces on which the IGP is to run.

The **network** command works for networks that the router learns dynamically or that are configured as static routes. Any routes that you inject into BGP via the **network** command must have an exact match with a corresponding IGP route. In the example below, network 2.2.2.0/24 must have a corresponding static route or an entry in the IGP routing table. The origin attribute of routes that are injected into BGP by means of the **network** command is set to IGP.

The following commands configure R3 to advertise network 2.2.2.0/24:

```
R3(config)# router bgp 230
R3(config-router)# neighbor 172.16.134.1 remote-as 100
R3(config-router)# network 2.2.2.0 mask 255.255.255.0
```

The **network** router configuration command causes R3 to generate an entry in the BGP table for network 2.2.2.0/24.

Summary

This section summarizes the key points discussed in this lesson.

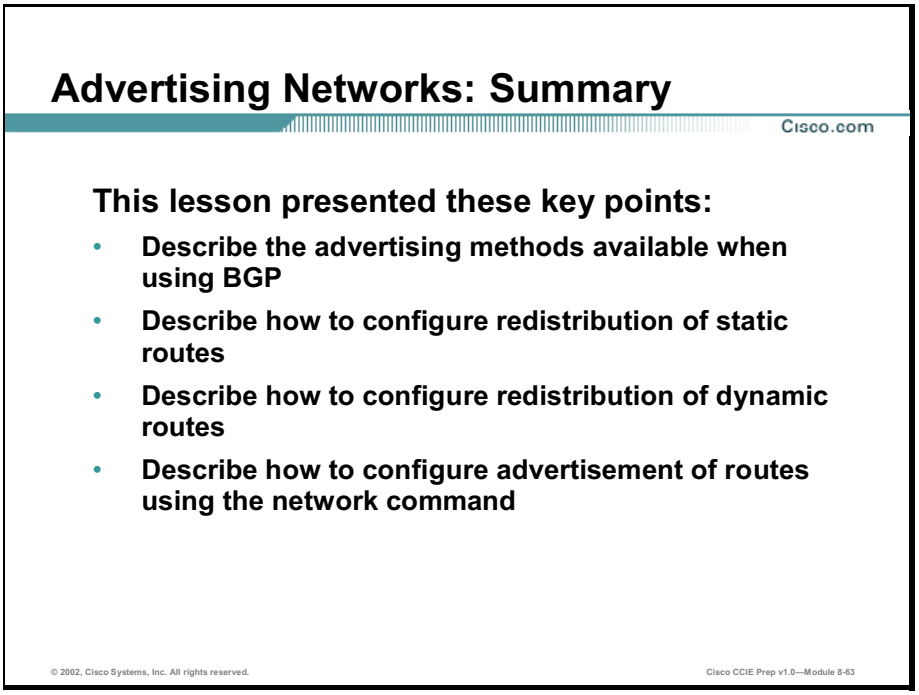

Next Steps

After completing this lesson, go to:

- BGP Advanced Options

References

For additional information, refer to these resources:

- *Internet Routing Architectures* by Sam Halabi

Lesson Assessment (Quiz)

Q1) Which of the following is **NOT** a valid method for advertising a route with Border Gateway Protocol (BGP)?

A) Redistributing static routes

B) Redistributing dynamic routes

C) Redistributing BGP into an Interior Gateway Protocol (IGP) such as Opent shortest Path First (OSPF)

D) Using the **network** command

Q2) Which of the following is usually discouraged?

A) Redistributing an IGP into BGP

B) Redistributing BGP into an IGP

C) Redistributing static routes that point to null 0

D) All of the above

Q3) When performing redistribution of any kind, which of the following commands is usually required?

A) ip route 0.0.0.0 0.0.0.0 <*ip-address*>

B) default-metric

C) ip classless

D) ip subnet-zero

Q4) Which of the following commands would you issue to redistribute EIGRP 10 into BGP Autonomous System (AS) 200?

A) R1(config-router)# redistribute eigrp 10

B) R1(config)# router eigrp 10

C) R1(config-router)# redistribute bgp 200

D) R1(config-router)# default-metric 1000 200 255 1 1500

Q5) Which command should be issued after modifying your configuration to implement redistribution?

A) default-metric

B) ip route <*ip-address*> <*mask*> null 0

C) clear ip bgp *

D) ip route 0.0.0.0 0.0.0.0 null 0

BGP Advanced Options

Overview

The Border Gateway Protocol (BGP) is a distance-vector routing protocol that is used to exchange routing information among different Autonomous Systems (AS). This lesson will discuss the more advanced topics of BGP, such as using private AS numbers, route dampening, route aggregation, and attribute modification.

Importance

BGP version 4 is the current exterior routing protocol used on the Internet. When connecting a multi-homed enterprise to the Internet, it is vital that the administrator know BGP properties, configurations and troubleshooting tips.

Objectives

Upon completing this lesson, you will be able to:

- Configure an AS using Private AS numbers
- Define and configure route dampening
- Define and configure route aggregation
- Perform conditional advertisement and route filtering
- Perform attribute modification
- Define peer groups and how they are configured

Learner Skills and Knowledge

To fully benefit from this lesson, you must have these prerequisite skills and knowledge:

- Passed the Cisco Certified Internetworking Expert (CCIE) written qualification exam
- Knowledge of Transmission Control Protocol/Internet Protocol (TCP/IP), access control lists, and Interior Gateway Protocols (IGPs)

Outline

This lesson includes these sections:

- Overview
- Using Private AS Numbers
- Dampening
- Route Aggregation
- Conditional Advertisement and Route Filtering
- Peer Groups
- Summary
- Lesson Assessment (Quiz)

Using Private AS Numbers

Since the Internet community has limited Autonomous System (AS) numbers, it is very difficult to obtain a "real" AS number.

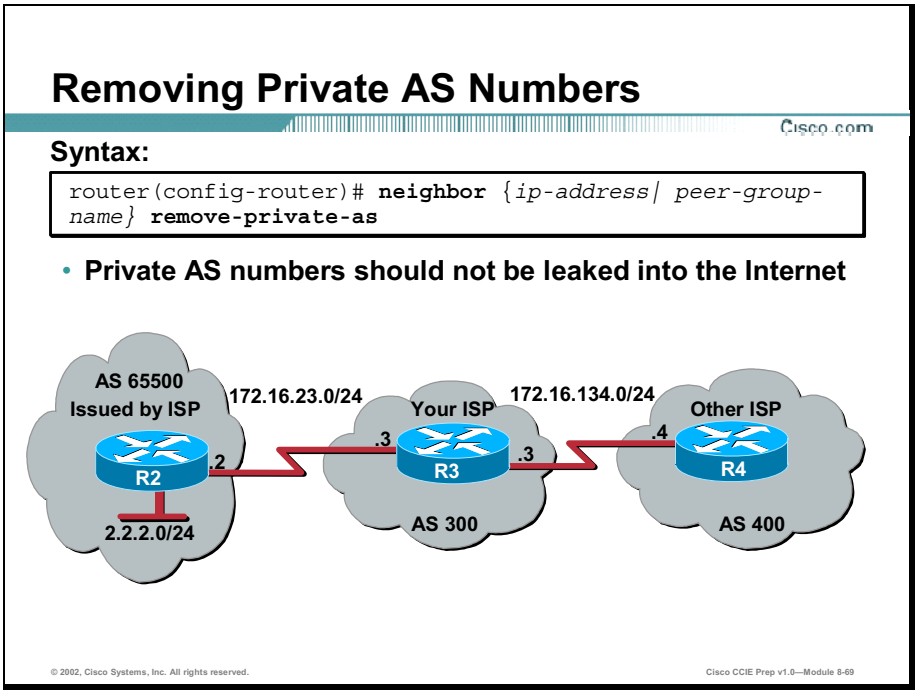

To overcome this limitation, the BGP spec specifies the use of private AS numbers, which range from 64152 to 65535. Your Internet Service Provider (ISP) can assign you a private AS, but that AS should not be advertised to the Internet community (other ISPs). To remove the private AS from updates, your ISP would issue the following command on peer statements to other ISPs.

neighbor {*ip-address | peer-group-name*} **remove-private-as**

Here is a simple scenario where R3 is your ISP's border router. The ISP has assigned you a private AS number of 65500, which is configured on R2. R2 is advertising network 2.2.2.0/24 to R3. In turn, R3 is advertising this network to R4, which is a different ISP.

When normal BGP peering is established you see the following in R4's BGP table.

```
R4# show ip bgp
BGP table version is 4, local router ID is 4.4.4.4
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
*> 2.2.2.0/24       172.16.134.3                           0 300 65500 i
```

Notice the path 300 65500. R3 should not advertise a private AS (65500) to R4. You need to add the following configuration command to R3.

```
R3(config)# router bgp 300

R3(config-router)# neighbor 172.16.134.4 remove-private-as
```

After clearing the BGP session to R4, view the BGP table on R4.

```
R4# show ip bgp
BGP table version is 6, local router ID is 4.4.4.4
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
*> 2.2.2.0/24       172.16.134.3                           0 300 I
```

The private AS has been removed from the AS path.

Dampening

This section covers dampening.

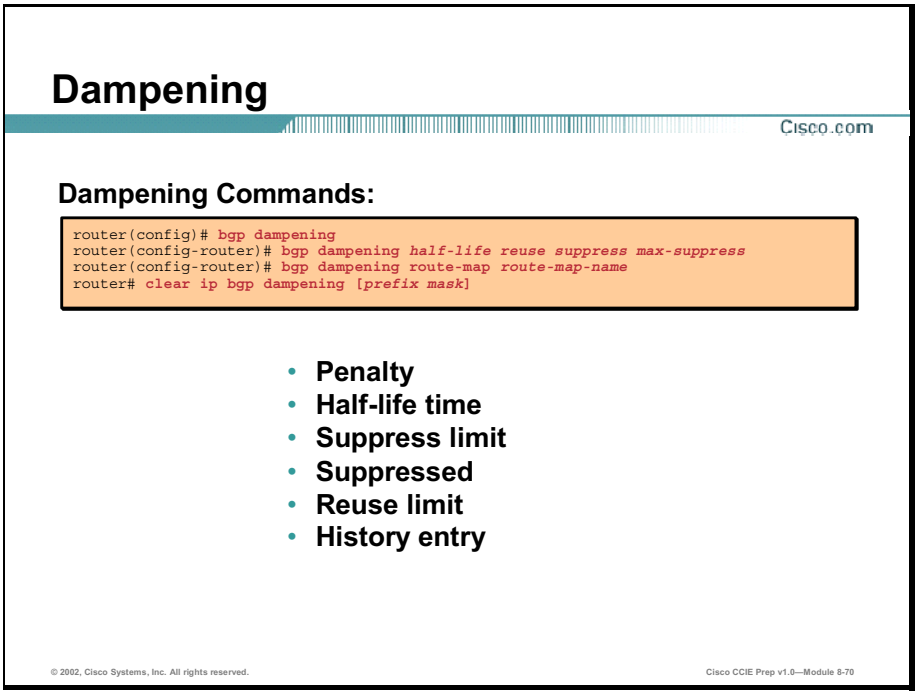

Border Gateway Protocol (BGP) sends a WITHDRAWN message to its peers when a prefix transitions from up to down. BGP sends an UPDATE message when the prefix transitions from down to up.

This is commonly referred to as a route flap, and can cause high CPU utilization while the BGP routes are converging. Also, if you are redistributing BGP into your Interior Gateway Protocol (IGP), this flapping can cause it to become less stable. Route flap dampening was introduced in Internetwork Operating System (IOS) Release 11.0 as a mechanism for minimizing the instability caused by route flapping. The following terms are used to describe route flap dampening:

- **Penalty:** A numeric value that is assigned to a route when it flaps. The default value is 1000.
- **Half-life time:** A configurable numeric value that describes the time required to reduce the penalty by one half. The default is 15 minutes.
- **Suppress limit:** A numeric value that is compared with the penalty. If the penalty is greater than the suppress limit, the route is suppressed. The default value is 2000.
- **Suppressed:** A route that is not advertised even though it is up. A route is suppressed if the penalty is more than the suppressed limit.

- **Reuse limit:** A configurable numeric value that is compared with the penalty. If the penalty is less than the reuse limit, a suppressed route that is up will no longer be suppressed. The default is 750.

- **History entry:** An entry that is used to store flap information about a route that is down.

A route that is flapping receives a penalty of 1000 for each flap. When the accumulated penalty reaches the suppress limit (2000), BGP suppresses advertisement of the route, even if the route is up. The accumulated penalty is decremented in half for each half-life time interval (15 minutes). When the accumulated penalty is less than the reuse limit, the route is advertised again, if it is still up.

Dampening is not applied to routes that are learned via internal BGP (iBGP). This restriction avoids forwarding loops and prevents iBGP peers from having a higher penalty for routes that are external to the AS.

The following commands are used when configuring route flap dampening:

```
bgp dampening
bgp dampening half-life reuse suppress max-suppress
bgp dampening route-map route-map-name
clear ip bgp dampening [prefix mask]
```

By default bgp dampening is disabled. The command **bgp dampening** enables it.

To modify the damping timers, issue the **bgp dampening** half-life reuse suppress max-suppress command.

To enable bgp dampening and apply different dampening parameters to different prefixes based on IP-address or AS-path information use the **bgp dampening route-map** command.

To clear dampening information for a specific or all dampened routes (unsuppress suppressed routes) issue the **clear ip bgp dampening** command.

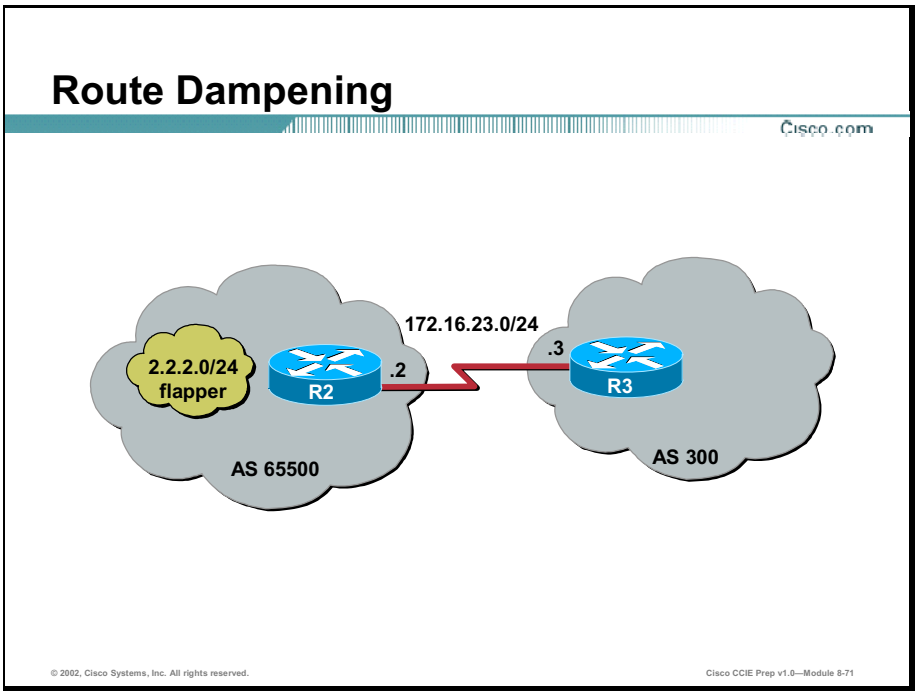

In this scenario, the advertised network 2.2.2.0/24 on R2 is flapping. R3 has been configured for route dampening as such:

```
R3(config)# router bgp 300
R3(config-router)# bgp dampening
```

The command **debug ip bgp dampening** has been issued and you receive the following output. During this time the network 2.2.2.0/24 has been flapping approximately every 30-45 seconds.

```
03:03:25: BGP(0): charge penalty for 2.2.2.0/24 path 65500 with halflife-time 15
 reuse/suppress 750/2000
03:03:25: BGP(0): flapped 1 times since 00:00:00. New penalty is 1000
R3#
03:04:18: BGP(0): charge penalty for 2.2.2.0/24 path 65500 with halflife-time 15
 reuse/suppress 750/2000
03:04:18: BGP(0): flapped 2 times since 00:00:53. New penalty is 1961
R3#
03:05:11: BGP(0): charge penalty for 2.2.2.0/24 path 65500 with halflife-time 15
 reuse/suppress 750/2000
03:05:11: BGP(0): flapped 3 times since 00:01:46. New penalty is 2886
R3#
03:06:16: BGP(0): suppress 2.2.2.0/24 path 65500 for 00:28:10 (penalty 2754)
03:06:16: halflife-time 15, reuse/suppress 750/2000
```

Now that the network 2.2.2.0/24 has been suppressed, you can view suppressed routes by issuing the following command:

```
R3# show ip bgp dampened-paths
BGP table version is 7, local router ID is 3.3.3.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? – incomplete

   Network          From             Reuse    Path
*d 2.2.2.0/24       172.16.23.2      00:27:40 65500 I
```

Notice that network 2.2.2.2 has been dampened and will not be used for 27 minutes and 40 seconds.

Route Aggregation

Border Gateway Protocol (BGP) allows the aggregation of specific routes into one route using the **aggregate-address** *address mask* command.

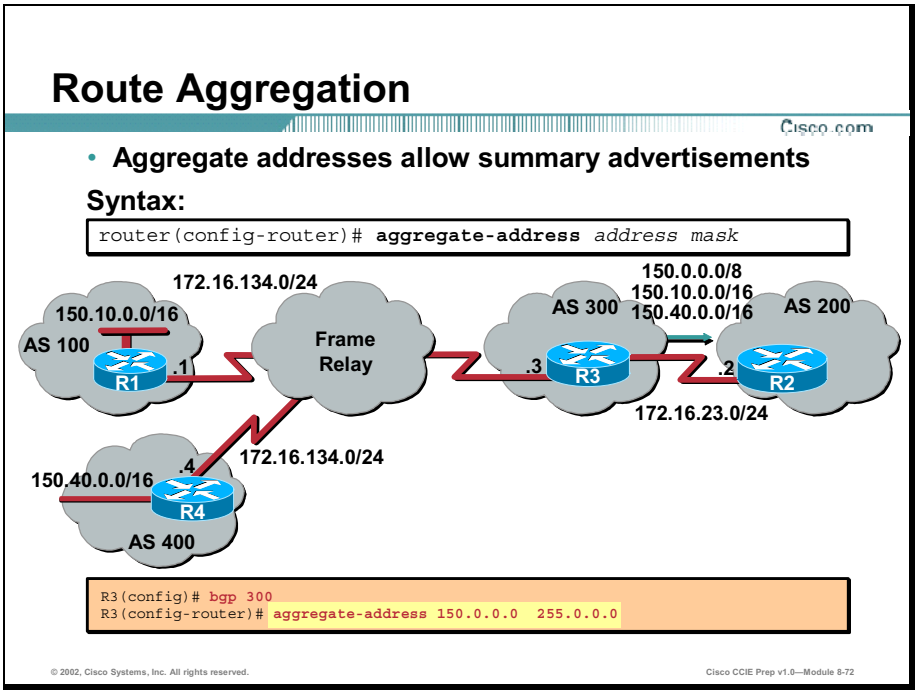

In the scenario, two different AS systems are sending class B networks to a 3rd AS.

- R1 is advertising the network 150.10.0.0/16 to R3.
- R4 is advertising the network 150.40.0.0/16 to R3.
- R3 has issued the following aggregate command in its BGP configuration.

```
R3(config)# router bgp 300
R3(config-router)# aggregate-address 150.0.0.0 255.0.0.0
```

If you look at the BGP table on R2, you see:

```
R2#show ip bgp
BGP table version is 4, local router ID is 10.10.10.10
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
*> 150.0.0.0/8      172.16.23.3                            0 300 i
```

```
*> 150.10.0.0       172.16.23.3                            0 300 100 i
*> 150.40.0.0       172.16.23.3                            0 300 400 I
```

R2 has received two class B networks:

```
The 150.10.0.0 network with the AS path set to 300 100.
The 150.40.0.0 network with the AS path set to 300 400.
```

You also have the aggregate created by R3. The 150.0.0.0/8 Classless Interdomain Routing (CIDR) network has the AS path set to 300.

AS path information is lost with the aggregate. When you configure aggregate-address without any arguments, it does not inherit the attributes (such as as_path or community) of the individual routes, which causes a loss of granularity. This may or may not be preferred, depending on the situation.

Notice that when R3 created the aggregate for the 150.0.0.0/8 network, the more specific routes were also sent and received by R2.

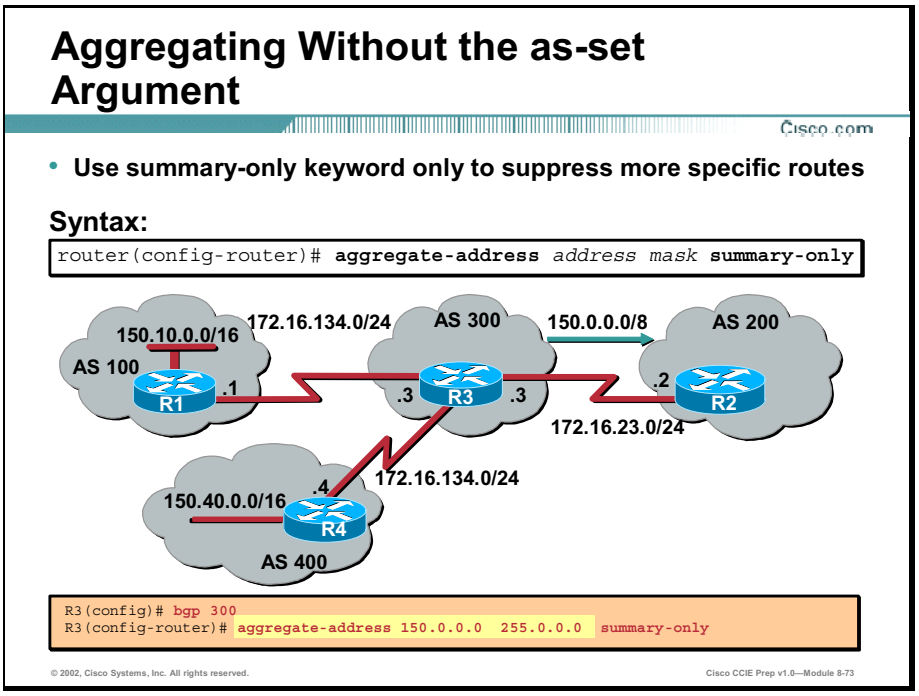

The form of the aggregate command that advertises the aggregate while suppressing all the more specific routes is shown.

aggregate-address *address address-mask* **summary-only**

Now look at the same scenario where two different AS systems are sending class B networks to a 3rd AS.

- R1 is advertising the network 150.10.0.0/16 to R3.
- R4 is advertising the network 150.40.0.0/16 to R3.
- R3 has issued the following aggregate command in its BGP configuration.

```
R3(config)# router bgp 300
R3(config-router)# aggregate-address 150.0.0.0 255.0.0.0 summary-only
```

The BGP table on R2 shows:

```
R2# show ip bgp
BGP table version is 6, local router ID is 10.10.10.10
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
*> 150.0.0.0/8      172.16.23.3                            0 300 i
```

Notice the more specific routes have been suppressed, but the problem of loss of information (AS path) has yet to be addressed.

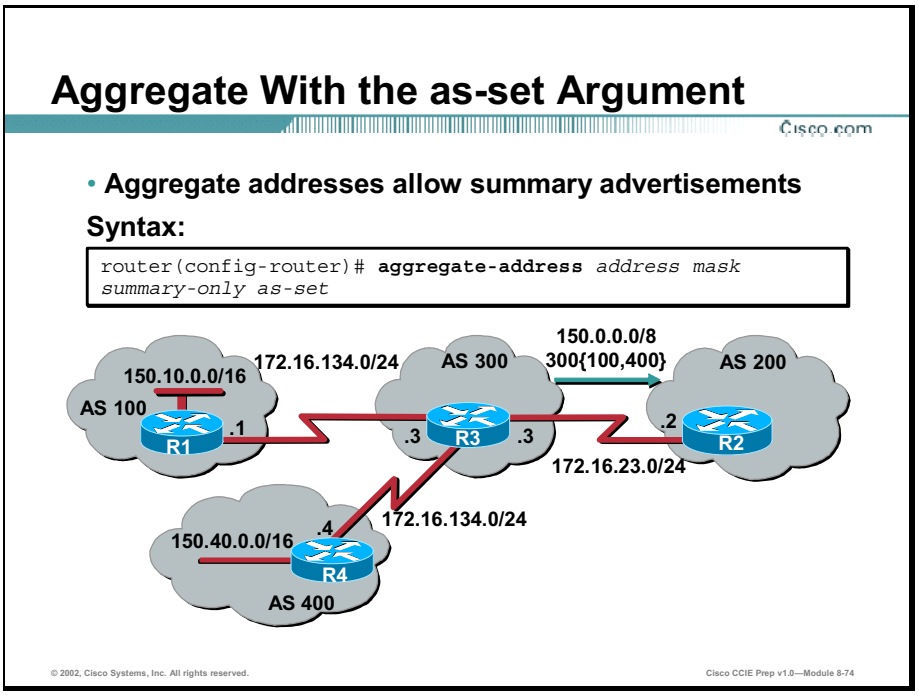

The form of the aggregate that advertises the aggregate while retaining the AS path information is shown. Using the **as-set** keyword creates an aggregate entry using the same rules that the command follows without this keyword, but the path advertised for this route will be an AS_SET consisting of all elements contained in all paths that are being summarized. Do not use this form of the **aggregate-address** command when aggregating many paths, because this route must be continually withdrawn and re-updated as AS path reachability information for the summarized routes changes.

aggregate-address *address address-mask* **as-set**

Here is the same scenario where two different AS systems are sending class B networks to a 3rd AS.

- R1 is advertising the network 150.10.0.0/16 to R3.
- R4 is advertising the network 150.40.0.0/16 to R3.
- R3 has issued the following aggregate command in its BGP configuration.

```
R3(config)# router bgp 300
R3(config-router)# aggregate-address 150.0.0.0 255.0.0.0 summary-only as-set
```

The BGP table on R2 shows:

```
R2# show ip bgp
BGP table version is 9, local router ID is 10.10.10.10
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
*> 150.0.0.0/8      172.16.23.3                           0 300 {100,400} I
```

You have received the summary, but now you have retained the AS path information from the more specific routes. Now, suppose there is a circumstance where you need to create an aggregate, but still need to allow some of the more specific routes to be propagated.

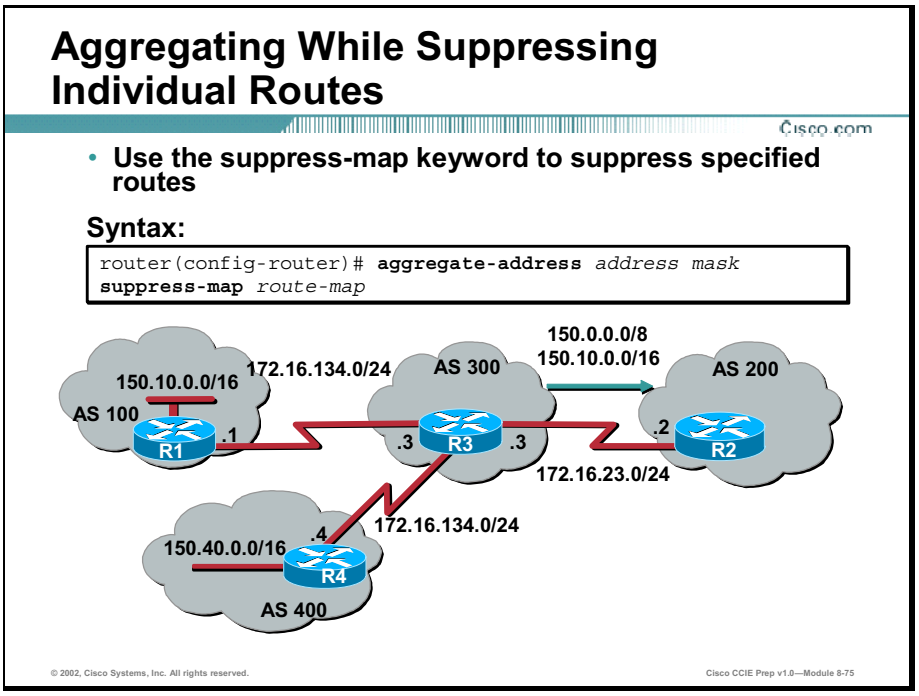

The form of the aggregate command that advertises the aggregate while suppressing only the more specific routes indicated by a route map is shown.

aggregate-address *address address-mask* **suppress-map** *route-map-name*

Look at the same scenario where two different AS systems are sending class B networks to a 3rd AS.

- R1 is advertising the network 150.10.0.0/16 to R3.
- R4 is advertising the network 150.40.0.0/16 to R3.
- R3 has issued the following commands in its BGP configuration.

```
R3(config)# router bgp 300
R3(config-router)# aggregate-address 150.0.0.0 255.0.0.0 suppress-map SUPPRESSR4
R3(config-router)# exit
R3(config)# route-map SUPPRESSR4 permit 10
R3(config-route-map)# match ip address 4
R3(config-route-map)# exit
R3(config)# access-list 4 permit 150.40.0.0 0.0.255.255
```

The BGP table on R2 shows:

```
R2#show ip bgp
BGP table version is 22, local router ID is 10.10.10.10
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
*> 150.0.0.0/8      172.16.23.3                            0 300 i
*> 150.10.0.0       172.16.23.3                            0 300 100 I
```

You have received the aggregate and have suppressed the R4 more specific route (150.40.0.0) while allowing all other more specific routes (150.10.0.0). This option can be very confusing to understand. In this case, any addresses associated with a permit statement in the access list will be denied. To make it easier to understand, read the suppress-map statement and access list as follows: You are permitting network 150.40.0.0/16 to be suppressed while all other routes will not be suppressed. In other words, the implicit deny all at the end of the access list denies all other routes from being suppressed. Any routes that are denied from being suppressed are allowed into the BGP table.

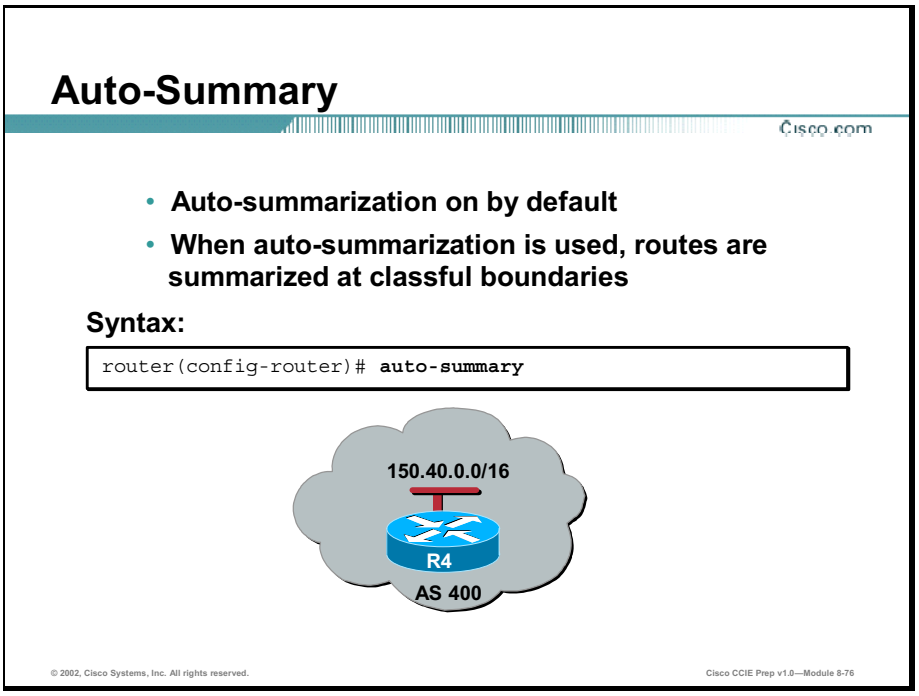

By default, BGP does not accept subnets redistributed from Interior Gateway Protocol (IGP). To advertise and carry subnet routes in BGP, use an explicit network command or the **no auto-summary** command. If you disable auto-summarization and have not entered a network command, you will not advertise network routes for networks with subnet routes unless they contain a summary route.

When you enable **auto-summary**, routes injected into BGP via redistribution are summarized at their classful boundary. **Auto-summary** does not apply to routes injected into BGP via the network command or through iBGP or external BGP (eBGP).

Auto-summary

Here is an example where R4 is redistributing static routes, connected routes, and routes learned via Enhanced Interior Gateway Routing Protocol (EIGRP) into BGP. Here is the R4 IP routing table.

```
R4# 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 not set
```

```
     1.0.0.0/24 is subnetted, 1 subnets
D       1.1.1.0 [90/2323456] via 172.16.70.3, 01:31:10, Ethernet0/0
     2.0.0.0/24 is subnetted, 1 subnets
D       2.2.2.0 [90/2323456] via 172.16.70.3, 01:31:10, Ethernet0/0
     3.0.0.0/24 is subnetted, 1 subnets
D       3.3.3.0 [90/409600] via 172.16.70.3, 01:31:10, Ethernet0/0
     4.0.0.0/24 is subnetted, 1 subnets
C       4.4.4.0 is directly connected, Loopback0
     5.0.0.0/24 is subnetted, 1 subnets
D       5.5.5.0 [90/2297856] via 172.16.45.5, 01:31:10, Serial0/1
     6.0.0.0/24 is subnetted, 1 subnets
D       6.6.6.0 [90/2809856] via 172.16.45.5, 01:31:11, Serial0/1
     172.16.0.0/16 is variably subnetted, 6 subnets, 2 masks
C       172.16.134.0/28 is directly connected, Serial0/0
D       172.16.56.0/24 [90/2681856] via 172.16.45.5, 01:31:11, Serial0/1
C       172.16.45.0/24 is directly connected, Serial0/1
D       172.16.23.0/24 [90/2195456] via 172.16.70.3, 01:31:11, Ethernet0/0
D       172.16.16.0/24 [90/2707456] via 172.16.70.3, 01:31:11, Ethernet0/0
C       172.16.70.0/24 is directly connected, Ethernet0/0
C    150.40.0.0/16 is directly connected, Loopback10
     30.0.0.0/24 is subnetted, 1 subnets
S       30.30.30.0 [1/0] via 150.40.0.2
```

Next, you redistribute EIGRP, connected, and static routes into R4's BGP. Here is R4's BGP table with the default of **auto-summary** enabled.

```
R4# show ip bgp
BGP table version is 32, local router ID is 4.4.4.4
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
*> 1.0.0.0          0.0.0.0                250         32768 ?
*> 2.0.0.0          0.0.0.0                250         32768 ?
*> 3.0.0.0          0.0.0.0                250         32768 ?
*> 4.0.0.0          0.0.0.0                  0         32768 ?
*> 5.0.0.0          0.0.0.0                250         32768 ?
*> 6.0.0.0          0.0.0.0                250         32768 ?
*> 30.0.0.0         0.0.0.0                250         32768 ?
*> 150.40.0.0       0.0.0.0                  0         32768 i
*> 172.16.0.0       0.0.0.0                  0         32768 ?
```

Notice that all routes have been summarized to their classful boundary. For instance, the 1.1.1.0/24 network in the IP routing table has been summarized to the classful network 1.0.0.0 in BGP.

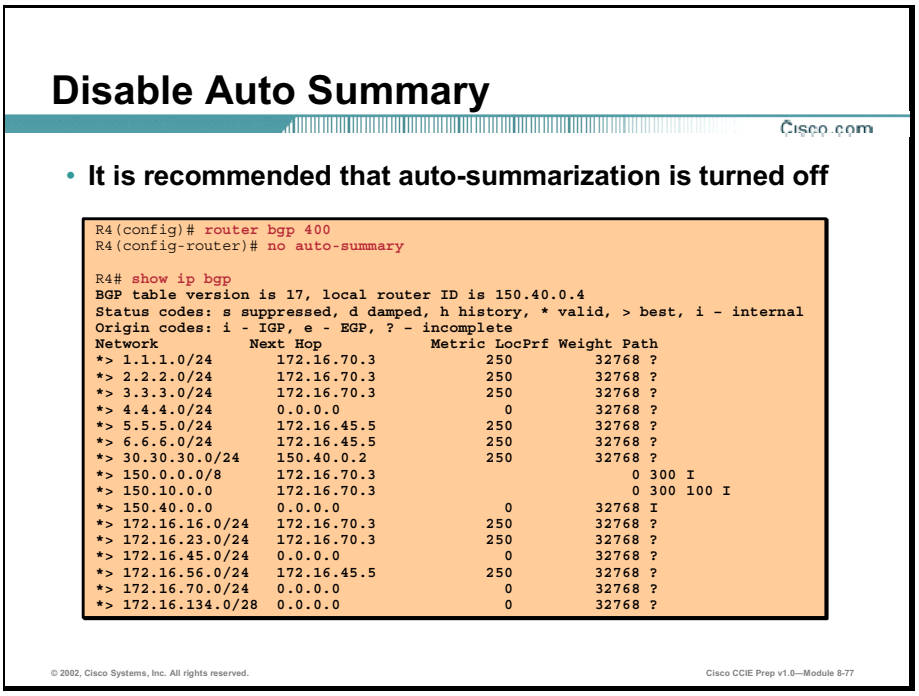

To disable auto-summary on R4, issue these commands:

```
R4(config)# router bgp 400
R4(config-router)# no auto-summary
```

Now, clear the BGP connections and then view R4's BGP routing table.

```
R4# clear ip bgp *
R4# show ip bgp
BGP table version is 17, local router ID is 150.40.0.4
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
*> 1.1.1.0/24       172.16.70.3            250         32768 ?
*> 2.2.2.0/24       172.16.70.3            250         32768 ?
*> 3.3.3.0/24       172.16.70.3            250         32768 ?
*> 4.4.4.0/24       0.0.0.0                  0         32768 ?
*> 5.5.5.0/24       172.16.45.5            250         32768 ?
*> 6.6.6.0/24       172.16.45.5            250         32768 ?
*> 30.30.30.0/24    150.40.0.2             250         32768 ?
*> 150.0.0.0/8      172.16.70.3                            0 300 i
*> 150.10.0.0       172.16.70.3                            0 300 100 i
*> 150.40.0.0       0.0.0.0                  0         32768 i
*> 172.16.16.0/24   172.16.70.3            250         32768 ?
```

```
*> 172.16.23.0/24   172.16.70.3           250         32768 ?
*> 172.16.45.0/24   0.0.0.0                0         32768 ?
*> 172.16.56.0/24   172.16.45.5           250         32768 ?
*> 172.16.70.0/24   0.0.0.0                0         32768 ?
*> 172.16.134.0/28  0.0.0.0                0         32768 ?
```

Conditional Advertisement and Route Filtering

The BGP conditional advertisement feature provides additional control of route advertisement depending on the existence of other prefixes in the BGP table.

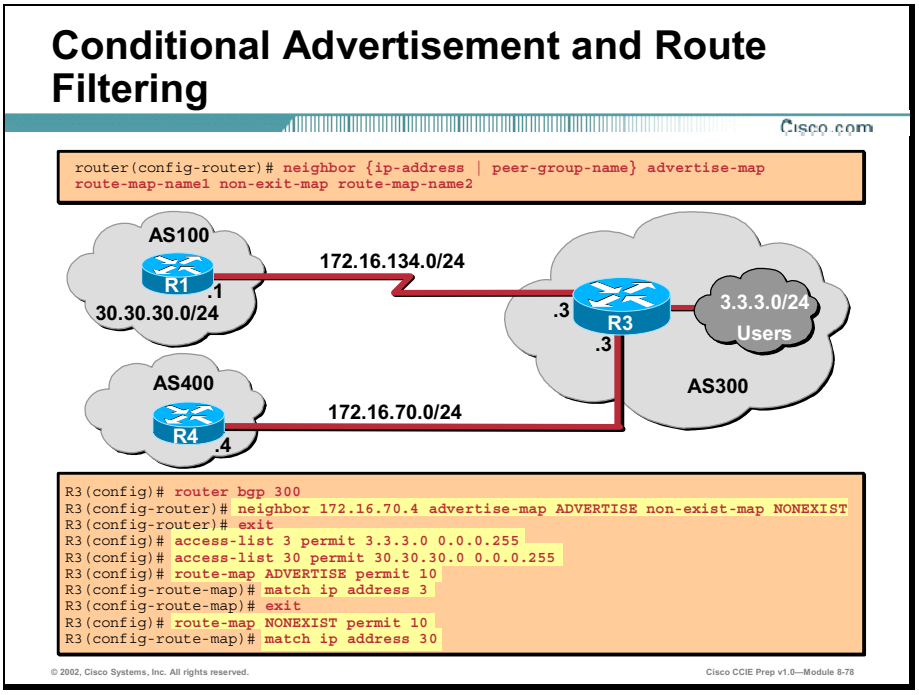

The BGP conditional advertisement feature provides additional control of route advertisement depending on the existence of other prefixes in the BGP table. Normally, routes are propagated regardless of the existence of a different path. The BGP conditional advertisement feature uses the **non-exist-map** and the **advertise-map** configuration commands to track routes by the route prefix. If a route prefix is not present in output of the **non-exist-map** command, then the route specified by the **advertise-map** command is announced. This feature is useful for multi-homed networks, in which some prefixes are advertised to one of the providers only if information from the other provider is missing (indicating a failure in the peering session or partial reachability).

neighbor {ip-address | peer-group-name} **advertise-map** *route-map-name1* **non-exit-map** *route-map-name2*

In the scenario, network users on network 3.3.3.0/24 in AS 300 have vital resources located in the 30.30.30.0/24 network so two-way connectivity must always be assured. The 30.30.30.0/24 network is being advertised to AS100. If all goes well AS 100 (our preferred path) will advertise the 30.30.30.0/24 network to AS300, which means you will advertise network 3.3.3.0/24 to AS100 to assure mutual connectivity.

If AS 300 loses connectivity to the 30.30.30.0/24 network, you want to use AS 400 as the preferred path and advertised the 3.3.3.0/24 network to it and withdraw the route to AS 100.

You verify connectivity by verifying that network 30.30.30.0/24 is indeed being advertised from AS 100. Your logic would proceed in the following way:

If AS 100 is advertising network 30.30.30.0/24 to AS 300

Then AS 300 will advertise network 3.3.3.0/24 to AS 100

Else

AS 300 will advertised network 3.3.3.0/24 to AS 400

Look at R3's BGP table before you perform conditional advertisement.

```
R3# show ip bgp
BGP table version is 4, local router ID is 3.3.3.3
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
*> 3.3.3.0/24       0.0.0.0                  0         32768 i
*> 4.4.4.0/24       172.16.70.4              0             0 400 i
*> 30.30.30.0/24    172.16.134.1             0             0 100 I
```

Here, you see the path to network 30.30.30.0/24 is indeed in our routing table sourced by R1.

And if you look at R1's BGP table you will see:

```
R1# show ip bgp
BGP table version is 20, local router ID is 1.1.1.1
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
*> 3.3.3.0/24       172.16.134.3             0             0 300 i
*> 4.4.4.0/24       172.16.134.3                           0 300 400 i
*> 30.30.30.0/24    0.0.0.0                  0         32768 I
```

Finally, look at R4's BGP table:

```
R4# show ip bgp
BGP table version is 5, local router ID is 4.4.4.4
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
*> 4.4.4.0/24       0.0.0.0                  0         32768 i
*> 30.30.30.0/24    172.16.70.3                            0 300 100 i
```

R3 will be configured in the following way:

```
R3(config)# router bgp 300
```



```
R3(config-router)# neighbor 172.16.70.4 advertise-map ADVERTISE non-exist-map
NONEXIST
R3(config-router)# exit
R3(config)# access-list 3 permit 3.3.3.0 0.0.0.255
R3(config)# access-list 30 permit 30.30.30.0 0.0.0.255
R3(config)# route-map ADVERTISE permit 10
R3(config-route-map)# match ip address 3
R3(config-route-map)# exit
R3(config)# route-map NONEXIST permit 10
R3(config-route-map)# match ip address 30
```

If R3 loses its route to 30.30.30.0/24, it will begin advertising network 3.3.3.0/24 to R4 as can be seen in the following output.

```
R4# show ip bgp
BGP table version is 5, local router ID is 4.4.4.4
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
*> 4.4.4.0/24       0.0.0.0                  0         32768 i
*> 3.3.3.0/24       172.16.70.3              0             0 300 100 i
```

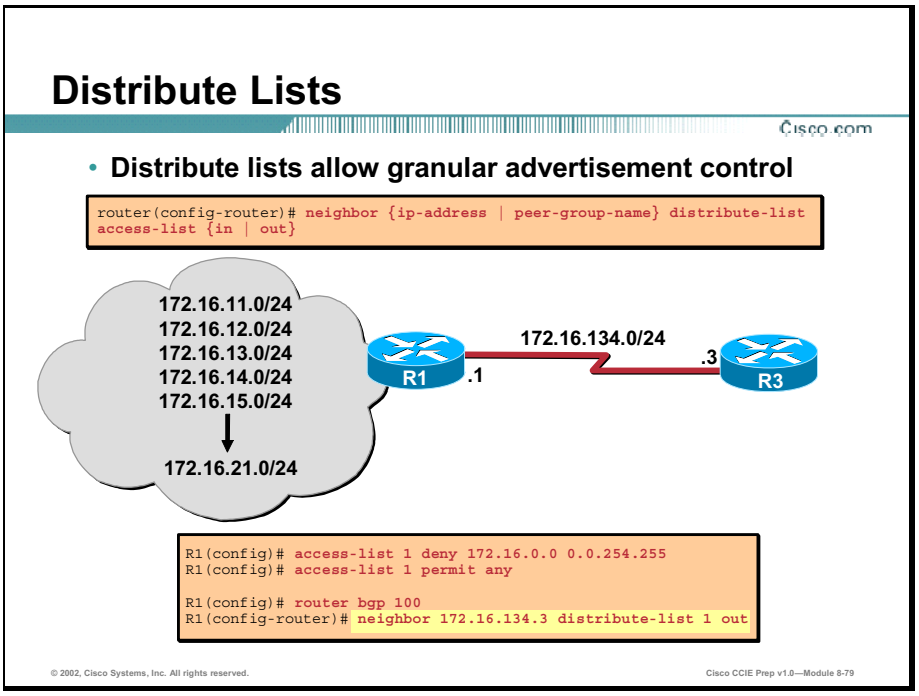

One method to filter BGP advertisements is to use distribute lists.

neighbor {*ip-address* | *peer-group-name*} **distribute-list** *access-list* {**in** | **out**}

Distribute lists are used in conjunction with access lists to filter routes in or out of BGP. Consider the following example.

R1 is advertising certain networks to R3. With no filtering enabled, here is R3's BGP table.

```
R3# sh ip bgp
BGP table version is 23, local router ID is 3.3.3.3
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
*> 3.3.3.0/24       0.0.0.0                  0         32768 I
*> 30.30.30.0/24    172.16.134.1             0             0 100 i
*> 172.16.11.0/24   172.16.134.1             0             0 100 i
*> 172.16.12.0/24   172.16.134.1             0             0 100 i
*> 172.16.13.0/24   172.16.134.1             0             0 100 i
*> 172.16.14.0/24   172.16.134.1             0             0 100 i
*> 172.16.15.0/24   172.16.134.1             0             0 100 i
*> 172.16.16.0/24   172.16.134.1             0             0 100 i
*> 172.16.17.0/24   172.16.134.1             0             0 100 i
*> 172.16.18.0/24   172.16.134.1             0             0 100 i
*> 172.16.19.0/24   172.16.134.1             0             0 100 i
```

```
*> 172.16.20.0/24   172.16.134.1             0             0 100 i
*> 172.16.21.0/24   172.16.134.1             0             0 100 i
```

To begin filtering, add the following statement to R1's neighbor statement to R3.

```
R1(config)# router bgp 300
R1(config-router)# neighbor 172.16.134.3 distribute-list 1 out
```

The command states that when R1 sends updates to R3; filter out networks according to access list 1.

In the first situation, you will only want to permit the odd 172.16.0.0 networks, as well as all other networks, to be advertised to AS 300. You can create access list 1 to perform this function.

```
R1(config)# access-list 1 deny   172.16.0.0 0.0.254.255
R1(config)# access-list 1 permit any
```

The first access list statement denies any even network in the 172.16.0.0 range.

```
access-list 1 deny   172.16.0.0 0.0.254.255
```

The second access list statement permits all other networks including the odd networks in the 172.16.0.0 range.

```
access-list 1 permit any
```

To verify first clear the BGP connections, then view R3's BGP table.

```
R3# show ip bgp
BGP table version is 72, local router ID is 3.3.3.3
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
*> 3.3.3.0/24       0.0.0.0                  0         32768 I
*> 30.30.30.0/24    172.16.134.1             0             0 100 i
*> 172.16.11.0/24   172.16.134.1             0             0 100 i
*> 172.16.13.0/24   172.16.134.1             0             0 100 i
*> 172.16.15.0/24   172.16.134.1             0             0 100 i
*> 172.16.17.0/24   172.16.134.1             0             0 100 i
*> 172.16.19.0/24   172.16.134.1             0             0 100 i
*> 172.16.21.0/24   172.16.134.1             0             0 100 I
```

In the second situation, you only want to permit the even 172.16.0.0 networks, as well as all other networks, to be advertised to AS 300. To do that, you modify the access list 1 as shown.

```
R1(config)# access-list 1 deny   172.16.1.0 0.0.254.255
R1(config)# access-list 1 permit any
```

The first access list statement denies any odd network in the 172.16.0.0 range.

```
access-list 1 deny   172.16.1.0 0.0.254.255
```

The second access list statement permits all other networks including the even networks in the 172.16.0.0 range.

```
access-list 1 permit any
```

To verify first clear the BGP connections, then view R3's BGP table.

```
R3# show ip bgp
BGP table version is 61, local router ID is 3.3.3.3
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
*> 3.3.3.0/24       0.0.0.0                  0         32768 I
*> 30.30.30.0/24    172.16.134.1             0             0 100 i
*> 172.16.12.0/24   172.16.134.1             0             0 100 i
*> 172.16.14.0/24   172.16.134.1             0             0 100 i
*> 172.16.16.0/24   172.16.134.1             0             0 100 i
*> 172.16.18.0/24   172.16.134.1             0             0 100 i
*> 172.16.20.0/24   172.16.134.1             0             0 100 I
```

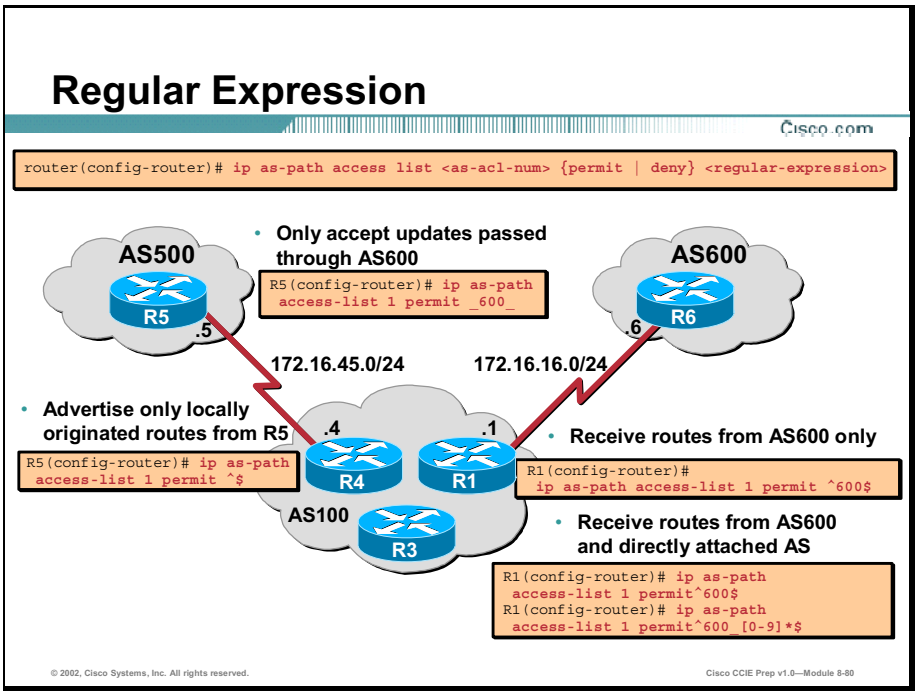

You can use regular expressions in the **ip as-path access-list** command with Border Gateway Protocol (BGP).

ip as-path access-list *<as-acl-num>* {**permit** | **deny**} <**regular-expression**>

Creating Regular Expressions

A regular expression can be a single-character pattern or a multiple-character pattern. That is, a regular expression can be a single character that matches the same single character in the input string, or multiple characters that match the same multiple characters in the input string. This section describes creating both single-character patterns and multiple-character patterns. It also discusses creating more complex regular expressions using multipliers, alternation, anchoring, and parentheses.

Single-Character Patterns

The simplest regular expression is a single character that matches itself in the input string. For example, the single-character regular expression 3 matches a corresponding 3 in the input string. You can use any letter (A-Z, a-z) or number (0-9) as a single-character pattern. The following examples are single-character regular expression patterns:

A

K

5

You can use a keyboard character other than a letter or a number—such as an exclamation point (!) or a tilde (~)—as a single-character pattern, but certain keyboard characters have

special meaning when used in regular expressions. The following table lists the keyboard characters with special meaning.

Table 7-1: Characters with Special Meaning Character

Characters with Special Meaning Character		Special Meaning
Period	.	Matches any single character, including white space
Asterisk	*	Matches 0 or more sequences of the pattern
Plus sign	+	Matches 1 or more sequences of the pattern
Question mark	?	Matches 0 or 1 occurrences of the pattern
Caret	^	Matches the beginning of the input string
Dollar sign	$	Matches the end of the input string
Underscore	_	Matches a comma (,), left brace ({), right brace (}), left parenthesis, right parenthesis, the beginning of the input string, the end of the input string, or a space
Brackets	[ ]	Designates a range of single-character patterns
Hyphen	-	Separates the end points of a range

To use these special characters as single-character patterns, remove the special meaning by preceding each character with a backslash (\). The following examples are single-character patterns matching a dollar sign, an underscore, and a plus sign, respectively:

\$

_

\+

You can specify a range of single-character patterns to match against a string. For example, you can create a regular expression that matches a string containing one of the following letters: a, e, i, o, and u. Only one of these characters must exist in the string for pattern matching to succeed. To specify a range of single-character patterns, enclose the single-character patterns in square brackets ([ ]). The order of characters within the brackets is not important. For example, [aeiou] matches any one of the five vowels of the lowercase alphabet, while [abcdABCD] matches any one of the first four letters of the lowercase or uppercase alphabet.

You can simplify ranges by entering only the end points of the range, separated by a dash (-).

Simplify the previous range as follows:

```
[a-dA-D]
```

To add a hyphen as a single-character pattern in your range, include another hyphen and precede it with a backslash:

```
[a-dA-D\-]
```


You can also include a right square bracket (]) as a single-character pattern in your range. To do so, enter the following:

```
[a-dA-D\-\]]
```

The previous example matches any one of the first four letters of the lowercase or uppercase alphabet, a hyphen, or a right square bracket.

You can reverse the matching of the range by including a caret (^) at the start of the range. The following example matches any letter except the ones listed:

```
[^a-dqsv]
```

The following example matches anything except a right square bracket (]) or the letter d:

```
[^\]d]
```

Multiple-Character Patterns

When creating regular expressions, you can also specify a pattern containing multiple characters. You create multiple-character regular expressions by joining letters, numbers, or keyboard characters that do not have special meaning. For example, a4% is a multiple-character regular expression. Precede keyboard characters that have special meaning with a backslash (\) when you want to remove their special meaning.

With multiple-character patterns, order is important. The regular expression a4% matches the character a followed by the number 4 followed by a percent (%) sign. If the input string does not have a4% in that order, pattern matching fails. The multiple-character regular expression a. uses the special meaning of the period character (.) to match the letter a followed by any single character. With this example, the strings ab, a!, or a2 are all valid matches for the regular expression.

You can remove the special meaning of the period character by preceding it with a backslash. In the expression **a\.**, only the string a. matches the regular expression.

Multipliers

You can create more complex regular expressions that instruct the Cisco IOS software to match multiple occurrences of a specified regular expression. To do so, you use some special characters with your single- and multiple-character patterns. The following table lists the special characters that specify "multiples" of a regular expression.

Table 7-2: Special Characters Used as Multipliers Character

Special Characters Used as Multipliers Character	Description
*	Matches 0 or more single- or multiple-character patterns
+	Matches 1 or more single- or multiple-character patterns
?	Matches 0 or 1 occurrences of the single or multiple-character pattern

The following example matches any number of occurrences of the letter a, including none:

a*

The following pattern requires that at least one letter a be present in the string to be matched:

a+

The following pattern matches the string bb or bab:

ba?b

The following string matches any number of asterisks (*):

**

To use multipliers with multiple-character patterns, enclose the pattern in parentheses. In the following example, the pattern matches any number of the multiple-character string ab:

(ab)*

As a more complex example, the following pattern matches one or more instances of alphanumeric pairs (but not none; that is, an empty string is not a match):

([A-Za-z][0-9])+

The order for matches using multipliers (*, +, or ?) is longest construct first. Nested constructs are matched from outside to inside. Concatenated constructs are matched beginning at the left side of the construct. Thus, the regular expression matches A9b3, but not 9Ab3 because the letter appears first in the construct.

Examples

Only Allow Networks Locally Originating from AS

For example, you want an AS to advertise only routes it locally originates. In another example, you want R5 to advertise only its locally originated routes. Apply the following outbound filter on R5.

```
ip as-path access-list 1 permit ^$
```

Only Allow Networks Originating from AS 600 to Enter R1

You want R1 to receive only the routes originated from AS 600 (and no Internet routes). You can apply an inbound access list on R1 as follows:

```
ip as-path access-list 1 permit ^600$
```

Only Allow Networks That Have Passed Through AS 600 to Enter AS 500

You want only the networks that have passed through AS 600 to enter AS 500 from R5. You can apply an inbound filter on R5.

```
ip as-path access-list 1 permit _600_
```

You can use an underscore (_) as the input string and output string in the **ip as-path access-list** command. Note that in this example, you do not use anchoring (for instance, there is no ^), so it doesn't matter what autonomous systems come before and after AS 600.

Only Allow Networks Originated from AS 600, and AS's Directly Attached to AS 600, to Enter R1

You want AS 100 to get networks originated from AS 600 and all directly attached AS's of AS 600. Apply the following inbound filter on R1.

```
ip as-path access-list 1 permit ^600$
ip as-path access-list 1 permit ^600_[0-9]*$
```

In the ip as-path access-list command, the carat (^) starts the input string and designates "AS". The underscore (_) means there is a null string in the string that follows "AS 600". The [0-9]* specifies that any connected AS with a valid AS number can pass the filter. The advantage of using the [0-9]* syntax is that it gives you the flexibility to add any number of AS's without modifying this command string.

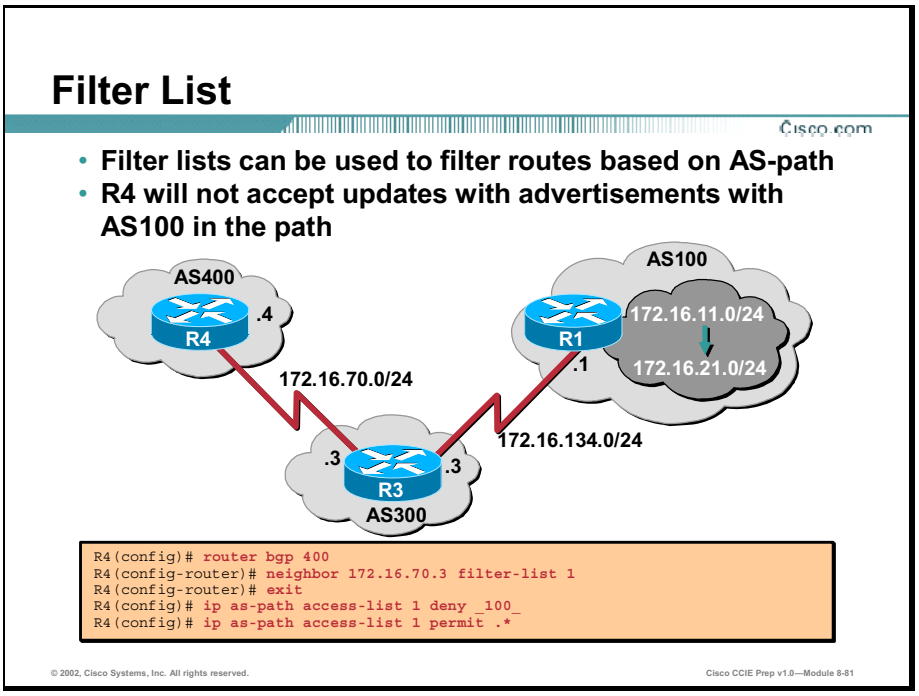

Filtering can also be performed via the AS path. To filter using the AS path, use the **neighbor filter-list** command. The complete syntax is shown.

neighbor {*ip-address* | *peer-group-name*} **filter-list** *as-path-list* {**in** | **out**}

Although **neighbor prefix-list** can be used as an alternative to the neighbor **distribute-list** command, do not use attempt to apply both **neighbor prefix list** and **neighbor distribute-list** filtering to the same neighbor. Only one filter list can be used per neighbor per direction.

In the above example R1 is advertising several networks to R3, which is in turn advertising those networks to R4 along with some of its own networks.

If you look at R4's BGP table before any filtering, you see:

```
R4# show ip bgp
BGP table version is 27, local router ID is 4.4.4.4
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
*> 3.3.2.0/24       172.16.70.3              0             0 300 i
*> 3.3.3.0/24       172.16.70.3              0             0 300 i
*> 3.3.4.0/24       172.16.70.3              0             0 300 i
*> 4.4.4.0/24       0.0.0.0                  0         32768 i
*> 30.30.30.0/24    172.16.70.3                            0 300 100 i
*> 172.16.11.0/24   172.16.70.3                            0 300 100 i
```

```
*> 172.16.12.0/24   172.16.70.3                          0 300 100 i
*> 172.16.13.0/24   172.16.70.3                          0 300 100 i
*> 172.16.14.0/24   172.16.70.3                          0 300 100 i
*> 172.16.15.0/24   172.16.70.3                          0 300 100 i
*> 172.16.16.0/24   172.16.70.3                          0 300 100 i
*> 172.16.17.0/24   172.16.70.3                          0 300 100 i
*> 172.16.18.0/24   172.16.70.3                          0 300 100 i
*> 172.16.19.0/24   172.16.70.3                          0 300 100 i
*> 172.16.20.0/24   172.16.70.3                          0 300 100 i
*> 172.16.21.0/24   172.16.70.3                          0 300 100 I
```

If at R4 you wanted to filter all networks with AS 100 in the path, you can use inbound AS path filtering.

```
R4(config)# router bgp 400
R4(config-router)# neighbor 172.16.70.3 filter-list 1
R4(config-router)# exit
R4(config)# ip as-path access-list 1 deny _100_
R4(config)# ip as-path access-list 1 permit .*
```

The **neighbor filter-list** command specifies that when R4 receives updates from neighbor 172.16.70.3 (R3), you should first pass the updates through AS path access list 1.

```
R4(config-router)# neighbor 172.16.70.3 filter-list 1
```

The first as-path access-list command denies any route that has AS 100 in the path.

```
R4(config)# ip as-path access-list 1 deny _100_
```

The second as-path access-list command permits all other routes.

```
R4(config)# ip as-path access-list 1 permit .*
```

After clearing the BGP connection and performing a show ip bgp, you will see the following.

```
R4# show ip bgp
BGP table version is 14, local router ID is 4.4.4.4
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
*> 3.3.2.0/24       172.16.70.3              0             0 300 i
*> 3.3.3.0/24       172.16.70.3              0             0 300 i
*> 3.3.4.0/24       172.16.70.3              0             0 300 i
```

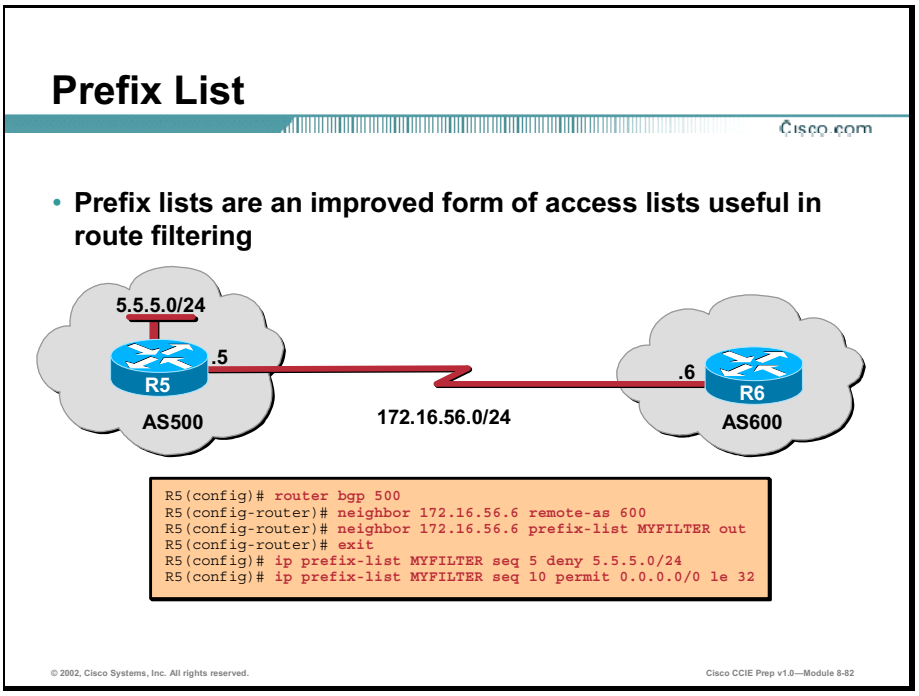

Another form of filtering is using prefix list filtering. Here, you can filter inbound or outbound routes based on the IP address and mask length. Only one prefix list can be used per neighbor, per direction. Using prefix lists are an alternative to using a distribution list with an extended access list. Two commands work in conjunction to perform prefix filtering:

neighbor {*ip-address | peer-group-name*} **prefix-list** *prefix-list-name* {**in** | **out**}

ip prefix-list <*prefix-list-num*> {**permit** | **deny**} <*ip_prefix*> [**ge** | **le**]] *network length*

There are two ways to block one or more networks from a Border Gateway Protocol (BGP) peer based on prefix. The first method uses the distribute-list out command and the second method uses the ip prefix-list command. The sample scenario will show the ip prefix-list method.

In this configuration, the **ip prefix-list** command matches any and denies the IP address range 5.5.5.0. Under the **router bgp 100** statement, specify the ip prefix-list command for the peer that you want.

The neighbor prefix-list command specifies you want to apply an outbound filter to updates directed to neighbor R6.

```
R5(config-router)# neighbor 172.16.56.6 remote-as 600
```

Prefix-list sequence 5 is denying the specific prefix 5.5.5.0/24.

```
R5(config)# ip prefix-list MYFILTER seq 5 deny 5.5.5.0/24
```

Prefix-list sequence 10 is permitting all other prefixes.

```
R5(config)# ip prefix-list MYFILTER seq 10 permit 0.0.0.0/0 le 32
```


Using Prefix Lists

Prefix lists can be used as an alternative to access lists in many BGP route-filtering commands.

The advantages of using prefix lists are:

- Significant performance improvement in loading and route lookup of large lists
- Support for incremental updates
  - Filtering using extended access lists does not support incremental updates.
- More user-friendly command-line interface
  - The command-line interface for using access lists to filter BGP updates is difficult to understand and use, since it uses the packet-filtering format.
- Greater flexibility
  - Before using a prefix list in a command, you must set up a prefix list, and you may want to assign sequence numbers to the entries in the prefix list.

How the System Filters Traffic by Prefix List

Filtering by prefix list involves matching the prefixes of routes with those listed in the prefix list. When there is a match, the route is used. The matching is similar to that of the access list. More specifically, whether a prefix is permitted or denied is based upon the following rules:

- An empty prefix list permits all prefixes.
- An implicit deny is assumed if a given prefix does not match any entries of a prefix list.
- When multiple entries of a prefix list match a given prefix, the sequence number of a prefix list entry identifies the entry with the lowest sequence number. In this case, the entry with the smallest sequence number is considered to be the "real" match.

The router begins the search at the top of the prefix list, with the sequence number 1. Once a match or deny occurs, the router does not need to go through the rest of the prefix list. For efficiency, you may want to place the most common matches or denies near the top of the list, using the argument *seq* in the **ip prefix-list** command. The **show** commands always include the sequence numbers in their output.

Sequence numbers are generated automatically unless you disable this automatic generation. If you disable the automatic generation of sequence numbers, you must specify the sequence number for each entry using the *seq-value* argument of the **ip prefix-list** command.

It does not matter if the default sequence numbers are used in configuring a prefix list, because a sequence number does not need to be specified when removing a configuration entry.

The optional keywords **ge** and **le** can be used to specify the range of the prefix length to be matched for prefixes that are more specific than *network*/*len*. An exact match is assumed when neither **ge** nor **le** is specified. The range is assumed to be from *ge-value* to 32 if only the **ge** attribute is specified, and from **len** to *le-value* if only the **le** attribute is specified.

A specified *ge-value* and/or *le-value* must satisfy the following condition:

```
len < ge-value <= le-value <= 32
```

For example, to deny all prefixes matching /24 in 128.0.0.0/8, you would use:

```
ip prefix-list abc deny  128.0.0.0/8 ge 24 le 24
```

Note You can specify sequence values for prefix list entries in any increments you want (the automatically generated numbers are incremented in units of 5). If you specify the sequence values in increments of 1, you cannot insert additional entries into the prefix list. If you choose very large increments, you could run out of sequence values.

To disable the automatic generation of sequence numbers, use the following command:

```
R5(config)# no ip prefix-list sequence-number
```

To delete a prefix list, use the following command in global configuration mode:

```
R5(config)# no ip prefix-list list-name
```

You can delete entries from a prefix list individually. To delete an entry in a prefix list, use the following command in global configuration mode:

```
R5(config)# no ip prefix-list seq seq-value
```

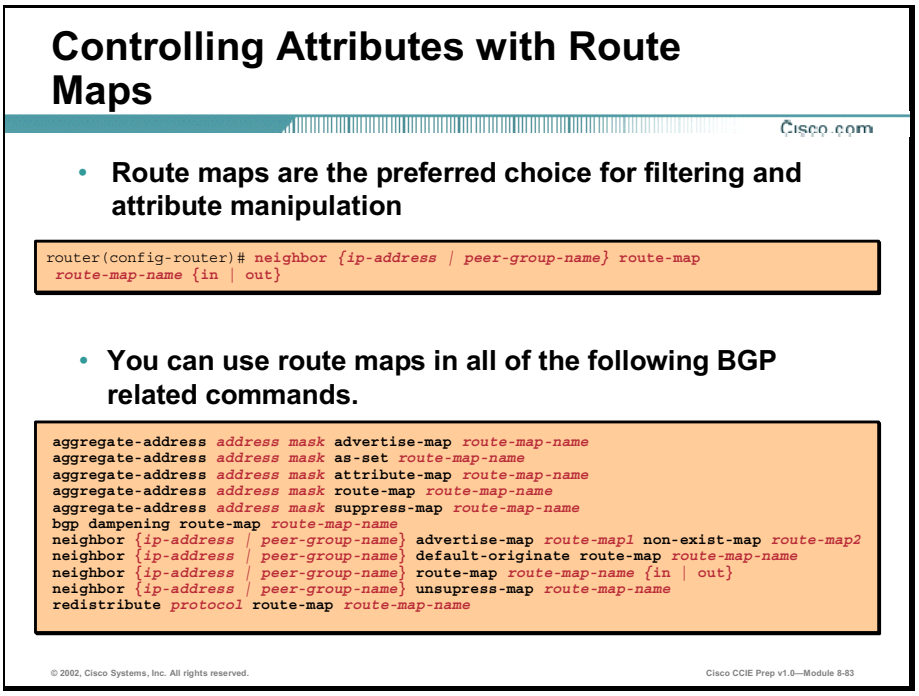

One of the most powerful tools in your arsenal is the route map. It is the tool of choice for route filtering as well as BGP attribute manipulation.

You can use a route map on a per-neighbor basis to filter updates and modify various attributes. A route map can be applied to either inbound or outbound updates. Only the routes that pass the route map are sent or accepted in updates.

Route Map Considerations

- Route map statements are numerically ordered and sequentially executed from lowest to highest number.
- An empty route map permit statement implicitly allows all routes
- An empty route map deny statement implicitly denies all routes
- Route map names are case sensitive

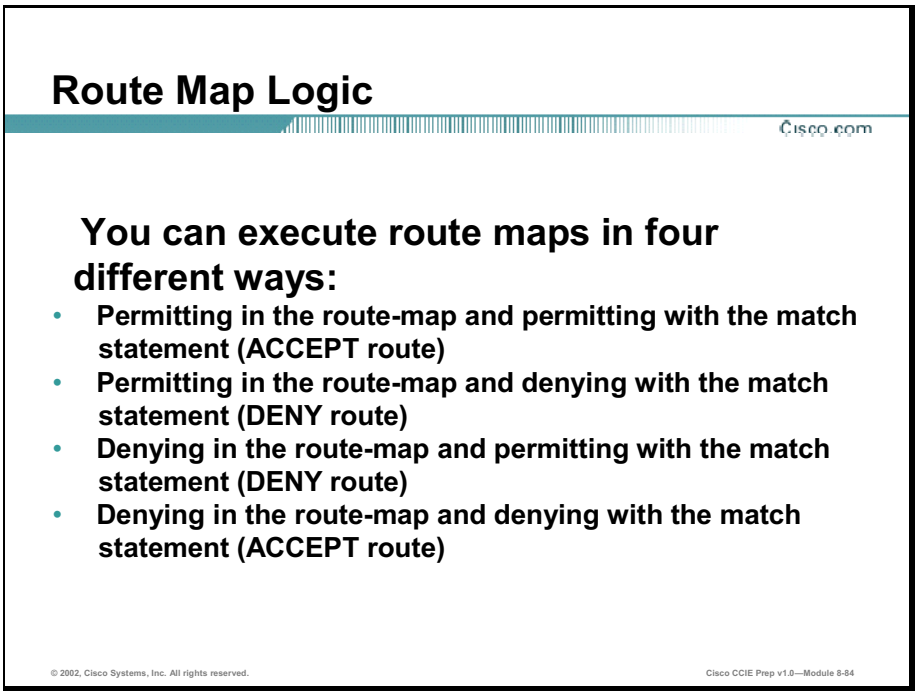

You need to understand in which conditions routes will be accepted, and if accepted, will the route-map execution terminate or will it continue processing the next statement.

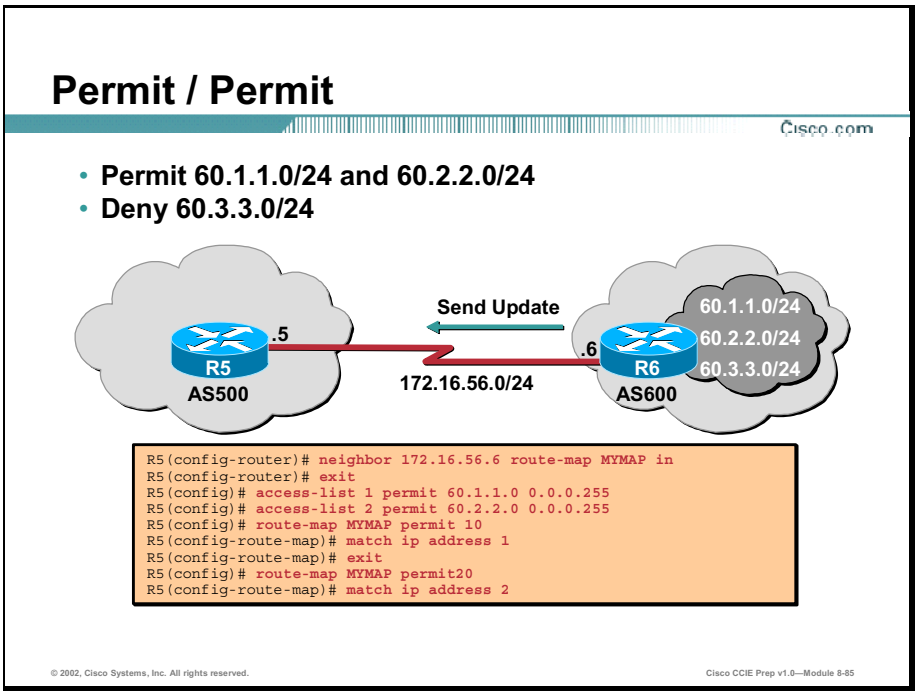

Permitting in the route-map and permitting with the conditional Access Control List (ACL) statement

With the permit/permit form the logic will follow this format:

If a match occurs

Then accept the route

set the attribute (if you using a set statement)

exit the route-map

Else execute the next route-map statement

Example:

```
R5(config-router)# neighbor 172.16.56.6 route-map MYMAP in
R5(config-router)# exit
R5(config)# access-list 1 permit 60.1.1.0 0.0.0.255
R5(config)# access-list 2 permit 60.2.2.0 0.0.0.255
R5(config)# route-map MYMAP permit 10
R5(config-route-map)# match ip address 1
R5(config-route-map)# exit
R5(config)# route-map MYMAP permit 20
R5(config-route-map)# match ip address 2
```

For each individual route:

- If route-map sequence number 10 is matched accept the route to 60.1.1.0/24, execute any set statements, and exit the route-map, do not continue processing.
- If route-map sequence number 10 is not matched, check sequence number 20.
- If route-map sequence number 20 is matched accept the route to 60.2.2.0/24, execute any set statements, and exit the route-map, do not continue processing.
- If route-map sequence number 20 is not matched, exit the route-map. Implicitly deny the route.

```
R5# show ip bgp
BGP table version is 3, local router ID is 5.5.5.5
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
*> 60.1.1.0/24      172.16.56.6              0             0 600 i
*> 60.2.2.0/24      172.16.56.6              0             0 600 i
```

You accepted networks 60.1.1.0/24 and 60.2.2.0/24. You denied all other networks (60.3.3.0/24). You could have included both networks in one access list if you had wanted to. By using separate access lists, you are able to apply different “set” parameters to each of the networks.

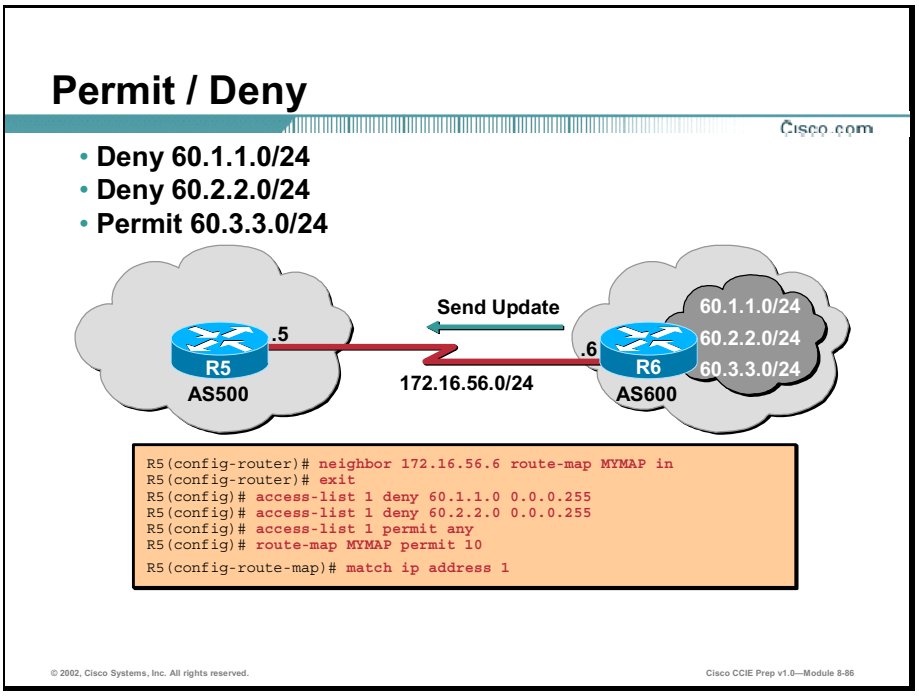

Permitting in the route-map and denying with the conditional (ACL) statement is the next example.

With the permit/deny form, the logic will follow this format:

If a match occurs

Then deny the route

exit the route-map

Example:

```
R5(config-router)# neighbor 172.16.56.6 route-map MYMAP in
R5(config-router)# exit
R5(config)# access-list 1 deny 60.1.1.0 0.0.0.255
R5(config)# access-list 1 deny 60.2.2.0 0.0.0.255
R5(config)# access-list 1 permit any
R5(config)# route-map MYMAP permit 10
R5(config-route-map)# match ip address 1
```

For each individual route:

- If route-map sequence number 10 is matched deny the route to 60.1.1.0/24 or 60.1.1.0/24 execute any set statements, and exit the route-map, do not continue processing. Without the permit any access list, all routes would be implicitly denied.

- If route-map sequence number 10 is not matched, exit the route-map. Implicitly deny the route.

```
R5# show ip bgp
BGP table version is 3, local router ID is 5.5.5.5
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
*> 60.3.3.0/24      172.16.56.6              0             0 600 i
```

You denied networks 60.1.1.0/24 and 60.2.2.0/24

You accepted all other networks (6.3.3.0/24)

This form is limited to only one route-map statement, because of the permit any in access list 1. Since all other routes would fall under this category the route-map sequence 10 would match and exit. You would never continue to any other route-map sequence number.

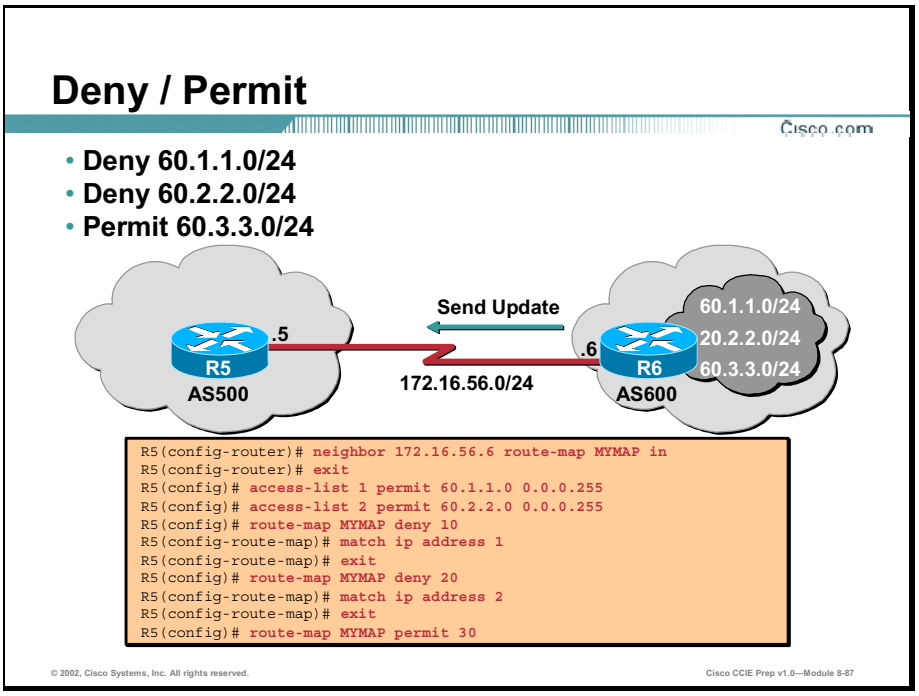

Denying in the route-map and permitting with the conditional (ACL) statement is shown.

With the deny/permit form, the logic will follow this format:

If a match occurs

Then deny the route

Exit the route-map

Example:

```
R5(config-router)# neighbor 172.16.56.6 route-map MYMAP in
R5(config-router)# exit
R5(config)# access-list 1 permit 60.1.1.0 0.0.0.255
R5(config)# access-list 2 permit 60.2.2.0 0.0.0.255
R5(config)# route-map MYMAP deny 10
R5(config-route-map)# match ip address 1
R5(config-route-map)# exit
R5(config)# route-map MYMAP deny 20
R5(config-route-map)# match ip address 2
R5(config-route-map)# exit
R5(config)# route-map MYMAP permit 30
```

For each individual route:

- If route-map sequence number 10 is matched deny, the route to 60.1.1.0/24, execute any set statements, and exit the route-map, do not continue processing.
- If route-map sequence number 10 is not matched, check sequence number 20.
- If route-map sequence number 20 is matched deny the route to 60.2.2.0/24, execute any set statements, and exit the route-map, do not continue processing.
- If route-map sequence number 20 is not matched, exit the route-map. Implicitly deny the route. This would deny all other routes sent from R6. Normally, you would like to receive all other routes, so route-map sequence number 30 was added to allow all other routes.

```
R5# show ip bgp
BGP table version is 2, local router ID is 5.5.5.5
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
*> 60.3.3.0/24      172.16.56.6              0             0 600 i

We denied networks 60.1.1.0/24 and 60.2.2.0/24.
We permitted all other routes (60.3.3.0/24)
```

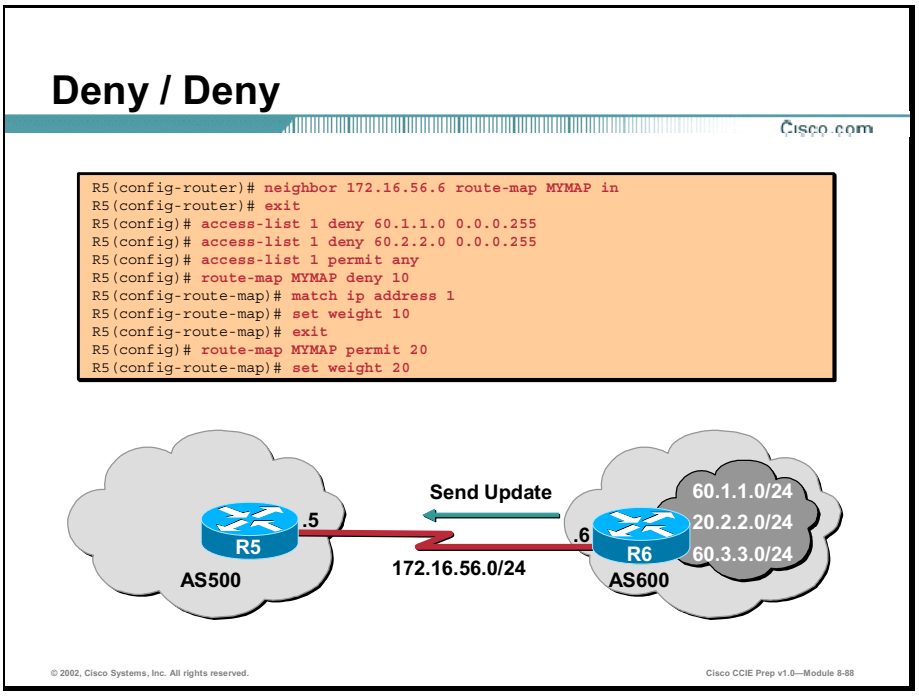

Denying in the route-map and denying with the conditional (ACL) statement is shown.

With the deny/deny form, the logic will follow the format:

If a match occurs

Then accept the route

execute the next route-map statement

The logic of this form is a little difficult to understand. To make it a little clearer, some set statements have been added.

```
R5(config-router)# neighbor 172.16.56.6 route-map MYMAP in
R5(config-router)# exit
R5(config)# access-list 1 deny 60.1.1.0 0.0.0.255
R5(config)# access-list 1 deny 60.2.2.0 0.0.0.255
R5(config)# access-list 1 permit any
R5(config)# route-map MYMAP deny 10
R5(config-route-map)# match ip address 1
R5(config-route-map)# set weight 10
R5(config-route-map)# exit
R5(config)# route-map MYMAP permit 20
R5(config-route-map)# set weight 20
```

For each individual route:

- If route-map sequence number 10 is matched accept the route (other than 60.1.1.0/24 or 60.2.2.0/24), execute any set statements, and exit the route-map, do not continue processing. In this case, accepting the route means to deny it.

- If route-map sequence number 10 is not matched (only networks 60.1.1.0/24 and 60.2.2.0/24), check sequence number 20.

- Sequence number 20 will implicitly permit all routes and execute any set statements. Only networks 60.1.1.0/24 and 60.2.2.0/24 will continue to sequence number 20.

```
R5# show ip bgp
BGP table version is 4, local router ID is 5.5.5.5
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
*> 60.1.1.0/24      172.16.56.6              0            20 600 i
*> 60.2.2.0/24      172.16.56.6              0            20 600 I
```

Notice that networks 60.1.1.0/24 and 60.2.2.0/24 have their weight set to 20, meaning they had passed route-map sequence number 20.

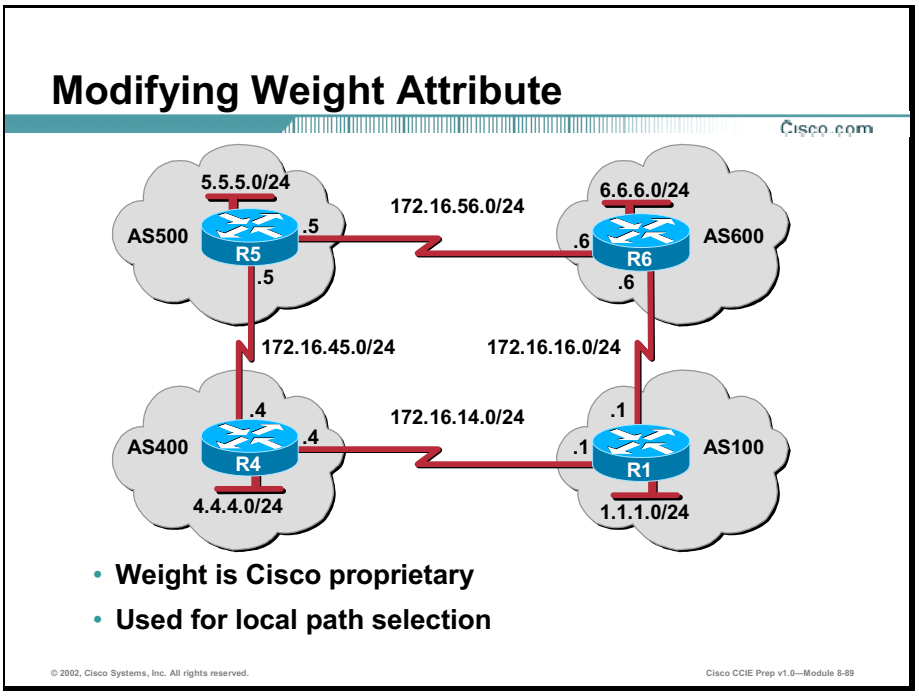

The recommended method to use when configuring route maps is the permit/permit method. Here, you can permit in the route map statement, then permit or deny as needed in the conditional statements (access lists, prefix lists, community lists, or AS path lists.) Here are some examples using this method:

Modifying Weight Attribute

- The weight attribute is a Cisco defined attribute.
- The weight is used for a best path selection process.
- The weight is assigned locally to the router.
- Weight is a value that only makes sense to the specific router and which is not propagated or carried through any of the route updates.
- A weight can be a number from 0 to 65535. Paths that the router originates have a weight of 32768 by default and other paths have a weight of zero.

The example shows how you can use route maps to modify incoming data from a neighbor.

In this scenario, you want to route through the AS 600 to reach network 1.1.1.0/24. To do that, create a route-map on R5 for routes received from AS 600. Any route received from 172.16.56.6 that matches the filter parameters set in access list 1 will have its weight set to 200, and it will be accepted. All other routes will be accepted and their weight will not be modified from the default value of 0. This will modify the weight attribute of the 1.1.1.0/24 network from R6, so the preferred route is through AS 600.

```
R5(config)# router bgp 500
R5(config-router)# neighbor 172.16.56.6 remote-as 600
R5(config-router)# neighbor 172.16.56.6 route-map MODWEIGHT in
R5(config-router)# exit
R5(config)# access-list 1 permit 1.1.1.0 0.0.0.255
R5(config)# route-map MODWEIGHT permit 10
R5(config-route-map)# match ip address 1
R5(config-route-map)# set weight 200
R5(config-route-map)# exit
R5(config)# route-map MODWEIGHT permit 20
```

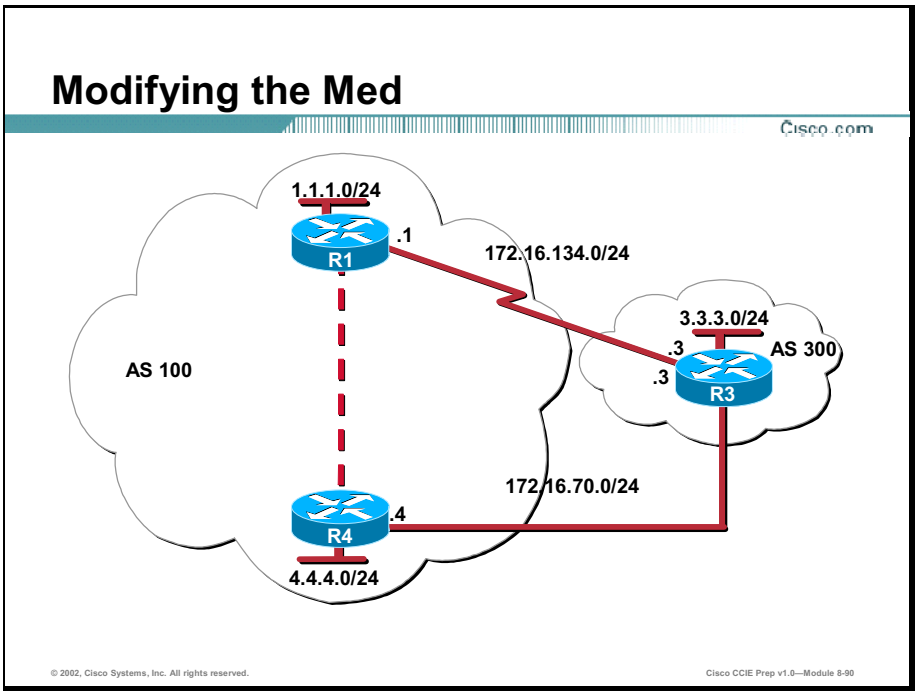

You can also modify the metric attribute.

- The metric attribute is also called Multi_Exit_Discriminator, MED (BGP4) or Inter-As (BGP3) is a hint to external neighbors about the preferred path into an AS.
- The metric attribute is a dynamic way to influence another AS on which way to choose in order to reach a certain route given that you have multiple entry points into that AS.
- A lower value of a metric is preferred.

Unlike local preference, metric is exchanged between (AS)s. A metric is carried into an AS but does not leave the AS. When an update enters the AS with a certain metric, that metric is used for decision making inside the AS. When the same update is passed on to a third AS, that metric will be set back to 0. The Metric default value is 0.

Unless otherwise specified, a router will compare metrics for paths from neighbors in the same AS. In order for the router to compare metrics from neighbors coming from different (AS)s the special configuration command **bgp always-compare-med** should be configured on the router.

In the example, AS 100 wants to influence AS 300 on the path to reach networks 1.1.1.0/24 and 4.4.4.0/24. To reach these networks, AS 300 should route to R4, not R1.

In the following example, MODMED will set the Multi Exit Discriminator (MED) to 1000 for the routes advertised from R4 and to 2000 from routes advertised from R1.

Before implementing the MED modification, look at R3's BGP table:

```
R3# show ip bgp
BGP table version is 4, local router ID is 3.3.3.3
```

```
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
*  1.1.1.0/24       172.16.70.4                            0 100 i
*>                  172.16.134.1             0             0 100 i
*> 3.3.3.0/24       0.0.0.0                  0         32768 i
*  4.4.4.0/24       172.16.70.4              0             0 100 i
*>                  172.16.134.1                           0 100 I
```

Here you see the preferred route to networks 1.1.1.0/24 and 4.4.4.0/24 is through R1 (172.16.134.1)

Now, implement the MED attribute modification on R4. Remember the route with the lowest MED will be the preferred route.

```
R4(config)# router bgp 100
R4(config-router)# neighbor 172.16.70.3 route-map MODMED out
R4(config-router)# exit
R4(config)# access-list 1 permit 1.1.1.0 0.0.0.255
R4(config)# access-list 1 permit 4.4.4.0 0.0.0.255
R4(config)# route-map MODMED permit 10
R4(config-route-map)# match ip address 1
R4(config-route-map)# set metric 1000
R4(config-route-map)# exit
R4(config)# route-map MODMED permit 20
```

```
R1(config)# access-list 1 permit 4.4.4.0 0.0.0.255
R1(config)# route-map MODMED permit 10
R1(config-route-map)# match ip address 1
R1(config-route-map)# set metric 2000
R1(config-route-map)# exit
R1(config)# route-map MODMED permit 20
```

After clearing the BGP connections on R3, issue the following command to verify the modifications:

```
R3# show ip bgp
BGP table version is 11, local router ID is 3.3.3.3
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
*> 1.1.1.0/24       172.16.70.4           1000             0 100 i
*                   172.16.134.1          2000             0 100 i
*> 3.3.3.0/24       0.0.0.0                  0         32768 i
*  4.4.4.0/24       172.16.134.1          2000             0 100 i
*>                  172.16.70.4           1000             0 100 I
```

R3 in AS 300 prefers the route through R4 to reach networks 1.1.1.0/24 and 4.4.4.0/24.

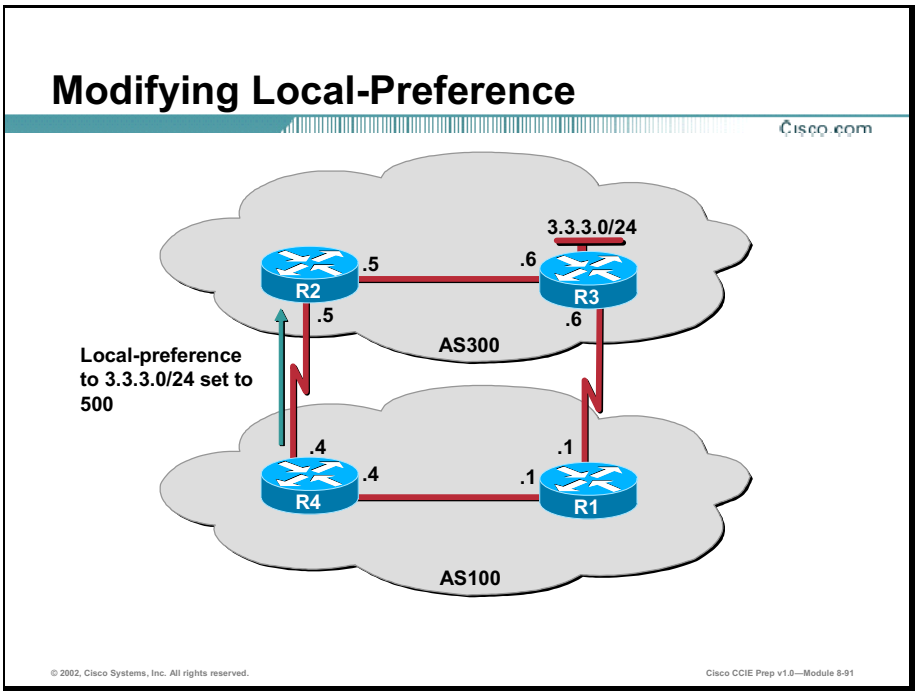

The following facts are part of modifying the local-preference:

- Local-preference is an indication to the AS about which path is preferred to exit the AS in order to reach a certain network.
- A path with a higher local-preference is more preferred.
- The default value for local-preference is 100.

Unlike the weight attribute that is only relevant to the local router, local-preference is an attribute that is exchanged among routers in the same AS.

Local-preference is set via the **bgp default local-preference** <*value*> command or with route-maps as will be demonstrated in the following example:

In this scenario, you want AS 100 to use the route between R4 and R3 to reach AS 300.

It is proper behavior to not accept any autonomous system path not matching the **match** clause of the route map. This means that you will not set the metric and the Cisco IOS software will not accept the route. However, you can configure the software to accept autonomous system paths not matched in the **match** clause of the route map command by using multiple maps of the same name, some without accompanying **set** commands.

If you view R1's BGP table, you see the following:

```
R1# show ip bgp
BGP table version is 6, local router ID is 1.1.1.1
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
*> 1.1.1.0/24       0.0.0.0                  0         32768 i
*> 3.3.3.0/24       172.16.134.3             0             0 300 i
* i                 172.16.70.3              0    100      0 300 i
*>i4.4.4.0/24       172.16.45.4              0    100      0 i
```

Here you see R1 is using its directly connected link to reach 3.3.3.0/24.

Next, configure R4 to modify its local preference for the 3.3.3.0/24 network.

```
R4(config)# router bgp 100
R4(config-router)# neighbor 172.16.0.3 route-map MYMAP in
R4(config-router)# exit
R4(config)# access-list 1 permit 3.3.3.0 0.0.0.255
R4(config)# route-map MYMAP permit 10
R4(config-route-map)# match ip address 1
R4(config-route-map)# set local-preference 500
R4(config-route-map)# exit
R4(config)# route-map MYMAP permit 20
```

```
R1# show ip bgp
BGP table version is 6, local router ID is 1.1.1.1
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
*> 1.1.1.0/24       0.0.0.0                  0         32768 i
*>i3.3.3.0/24       172.16.70.3              0    500      0 300 i
*                   172.16.134.3             0             0 300 i
*>i4.4.4.0/24       172.16.45.4              0    100      0 i
```

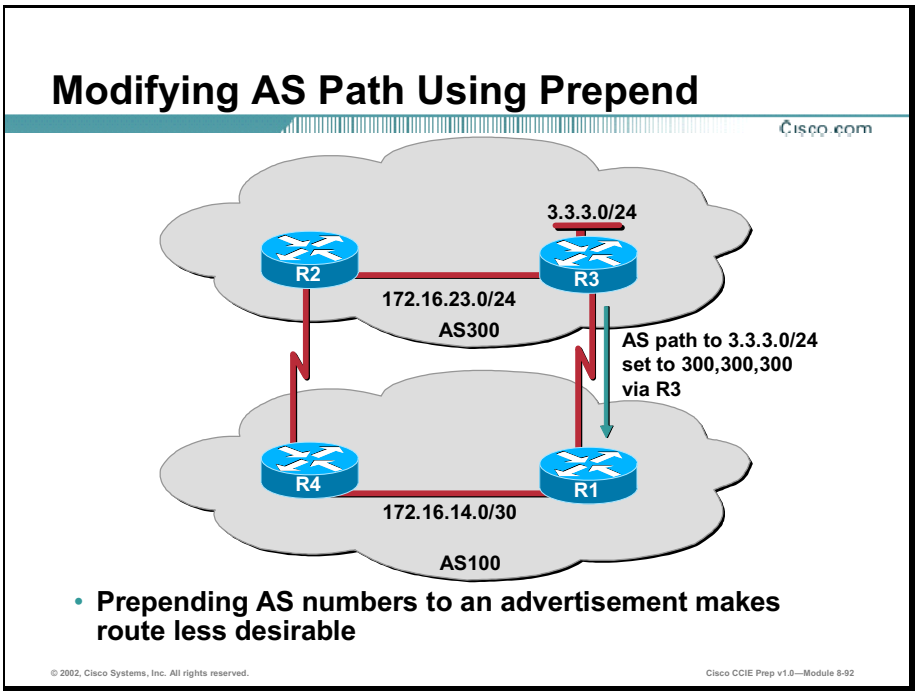

Another way to influence the path to a network is to modify the AS path. When you prepend AS paths to a prefix, you are making the route look less attractive.

In this scenario, you want to influence AS 100 on how it can reach the 3.3.3.0/24 network. You want to make sure that packets traveling from AS 100 to the 3.3.3.0/24 network travel over the link between R4 and R3.

The following example shows how the route map called **set-as-path** is applied to outbound updates to the neighbor 172.16.134.1. The route map will prepend the autonomous system path "300 300" to routes that pass access list 1. The second part of the route map is to permit the advertisement of other routes.

Before you begin configuration, look at the BGP table on R1.

```
R1# show ip bgp
BGP table version is 8, local router ID is 1.1.1.1
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
*> 1.1.1.0/24       0.0.0.0                  0         32768 i
* i3.3.3.0/24       172.16.23.2                   100      0 300 i
*>                  172.16.134.3             0             0 300 i
*>i4.4.4.0/24       172.16.45.4              0    100      0 i
```

Here you prefer the link between R1 and R3 to reach the 3.3.3.0/24 network. Lets perform the configuration on R3.

```
R3(config)# router bgp 300
R3(config-router)# network 3.3.3.0 mask 255.255.255.0
R3(config-router)# neighbor 172.16.70.4 remote-as 100
R3(config-router)# neighbor 172.16.134.1 remote-as 100
R3(config-router)# neighbor 172.16.134.1 route-map SET-AS-PATH out
R3(config-router)# exit
R3(config)# access-list 1 permit 3.3.3.0 0.0.0.255
R3(config)# route-map SET-AS-PATH permit 10
R3(config-route-map)# match address 1
R3(config-route-map)# set as-path prepend 300 300
R3(config-route-map)# exit
R3(config)# route-map SET-AS-PATH permit 20
```

If you view the BGP table on R1, you see the following:

```
R1# show ip bgp
BGP table version is 4, local router ID is 1.1.1.1
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
*> 1.1.1.0/24       0.0.0.0                  0         32768 i
*>i3.3.3.0/24       172.16.23.2                   100      0 300 i
*                   172.16.134.3             0            0 300 300 300 i
*>i4.4.4.0/24       172.16.45.4              0    100     0 i
```

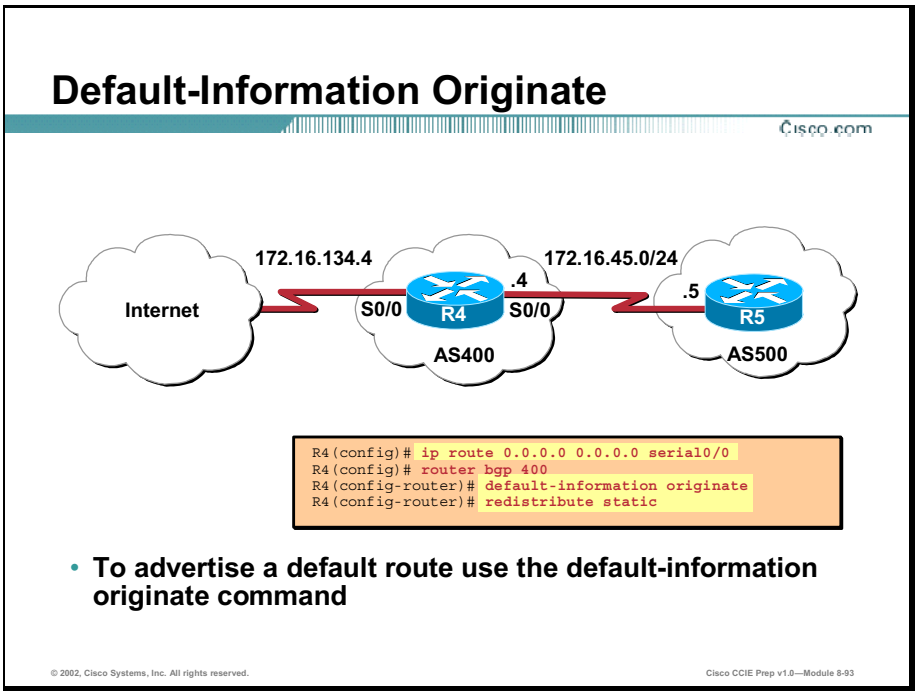

When you have a BGP speaker advertise a default route, you should issue the following command in router configuration mode:

```
Default-information originate
```

For example:

On R4 you issue the following command:

```
R4(config)# router bgp 400
R4(config-route)# default-information originate
```

After clearing the BGP connections, you look at R5 to see the following:

```
R5# show ip bgp
BGP table version is 17, local router ID is 5.5.5.5
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
*> 4.4.4.0/24       172.16.45.4              0             0 400 i
*> 5.5.5.0/24       0.0.0.0                  0         32768 i
```

You do not have a default network in the BGP table. This is because even though you have issued the proper command, R4 itself has no default network. If R4 has no default network, it will not advertise one. To remedy this issue the following command on R4:

```
R4(config)# ip route 0.0.0.0 0.0.0.0 serial0/0
```

After clearing the BGP connection, look at R5's BGP table once again.

```
R5# show ip bgp
BGP table version is 17, local router ID is 5.5.5.5
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
*> 4.4.4.0/24       172.16.45.4              0             0 400 i
*> 5.5.5.0/24       0.0.0.0                  0         32768 I
```

Again, you see that you have not received the default network on R5. This is because BGP requires not only the default-information **originate** command and a default-network created, it also requires that you redistribute the default-network into BGP. You need to issue this command on R4:

```
R4(config)# router bgp 400
R4(config-router)# redistribute static\
```

After clearing the BGP connection once more, look at R5's BGP table:

```
R5# show ip bgp
BGP table version is 17, local router ID is 5.5.5.5
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
*> 0.0.0.0          172.16.45.4              0             0 400 ?
*> 4.4.4.0/24       172.16.45.4              0             0 400 i
*> 5.5.5.0/24       0.0.0.0                  0         32768 i
```

Finally, you have received your default network on R5. Remember, when advertising a default network from BGP three items need to be completed on the router advertising the default network:

- Create a static default route (0.0.0.0 0.0.0.0)
- Redistribute the static route into BGP (redistribute static)
- Issue the **default-information originate** command

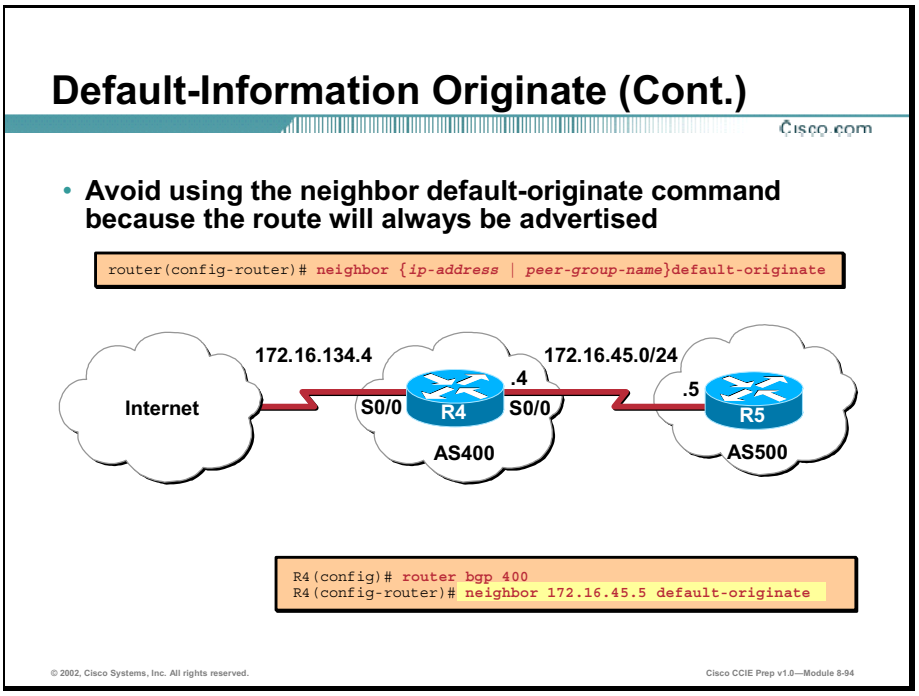

Another way of advertising a default route is via the neighbor default-originate command. This method is not recommended, as the advertising router will always advertise the default route, even if the route to the default network is down, or the router does not have a default route.

neighbor {*ip-address* | *peer-group-name*} **default-originate**

For example, R4 does not have a default route in its IP routing table, yet it will always advertise itself as being the default network for R5 when the following command is issued on R4:

```
R4(config)# router bgp 400
R4(config-router)# neighbor 172.16.45.5 default-originate
```

You can verify by performing a show ip bgp on R5:

```
R5# show ip bgp
BGP table version is 17, local router ID is 5.5.5.5
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
*> 0.0.0.0          172.16.45.4              0             0 400 ?
*> 4.4.4.0/24       172.16.45.4              0             0 400 i
*> 5.5.5.0/24       0.0.0.0                  0         32768 i
```

Peer Groups

Using peer groups can significantly reduce Central Processing Unit (CPU) calculations if you have many peers with the same update policies. Without peer groups, BGP would have to calculate an update policy for each neighbor even though the updates are identical.

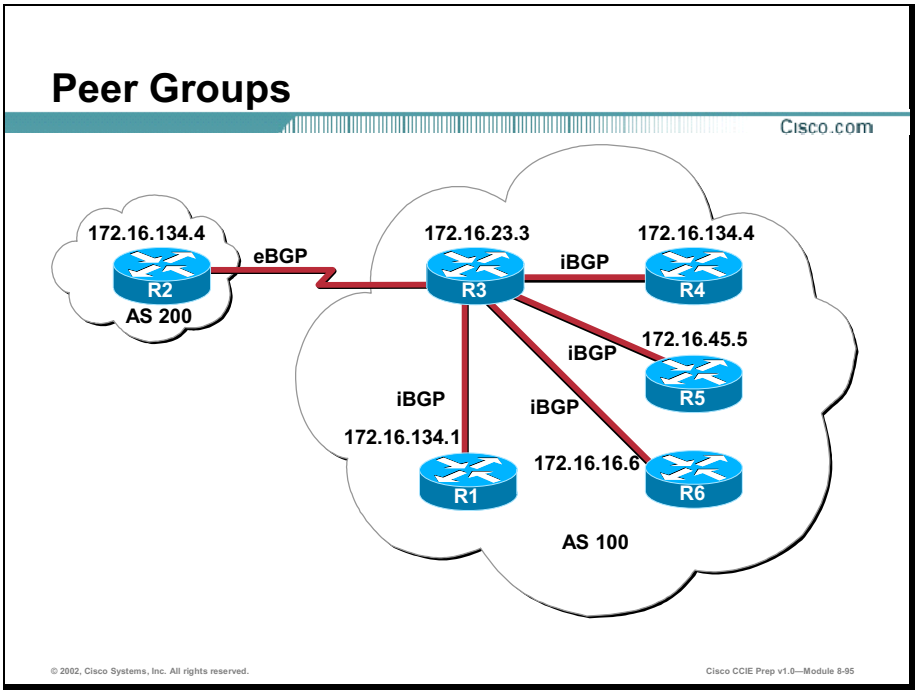

A BGP peer group is a group of BGP neighbors that share the same update policies. Update policies are usually set by route maps, distribution lists, and filter lists. Instead of defining the same policies for each individual neighbor, you define a peer group name and assign policies to the peer group.

Members of a peer group inherit all of the configuration options of the peer group. Peer group members can also be configured to override configuration options if the options do not affect outgoing updates. That is, you can only override options that are set for incoming updates.

The commands used to create peer groups and place neighbors in those groups are listed below.

neighbor *peer-group-name* **peer-group**
neighbor *ip-address* **peer-group** *peer-group-name*

The **neighbor** *peer-group-name* **peer-group** command creates a BGP peer group.

The **neighbor** *ip-address* **peer-group** *peer-group-name* command adds a neighbor to an existing peer group.

Consider a scenario where update policies to multiple neighbors can be simplified by using peer groups.

In this scenario Routers R1, R4, R5, and R6 have the exact same update policies.

- They are each part of AS 100
- They have the same outbound route map (INTERNAL)
- They apply the same outbound filter rules (filter-list 1)
- They apply the same inbound filter rules (filter-list 2)
- They each require the next hop modified
- They all allow inbound soft-reconfiguration
- Each router uses its own loopback 0.

The only difference is with R1, which needs an additional inbound filter (filter-list 3).

Normally, each router would require a neighbor statement for each of the above policies. In this scenario, you have 7 policies applied to 4 routers. That equates to 28 policy statements. You can significantly reduce those numbers by creating a single policy placed in a peer group named IBGPPEERS. Then let each neighbor know they are part of that peer group. This is an example of what the neighbor peer statements would look like.

```
R3(config)# router bgp 100
R3(config-router)# neighbor IBGPPEERS peer-group
R3(config-router)# neighbor IBGPPEERS remote-as 100
R3(config-router)# neighbor IBGPPEERS route-map INTERNAL out
R3(config-router)# neighbor IBGPPEERS filter-list 1 out
R3(config-router)# neighbor IBGPPEERS filter-list 2 in
R3(config-router)# neighbor IBGPPEERS next-hop-self
R3(config-router)# neighbor IBGPPEERS soft-reconfiguration in
R3(config-router)# neighbor IBGPPEERS update-source loopback 0
R3(config-router)# neighbor 4.4.4.4 peer-group IBGPPEERS
R3(config-router)# neighbor 5.5.5.5 peer-group IBGPPEERS
R3(config-router)# neighbor 6.6.6.6 peer-group IBGPPEERS
R3(config-router)# neighbor 1.1.1.1 peer-group IBGPPEERS
R3(config-router)# neighbor 1.1.1.1 filter-list 3 in
```

A peer-group named IBGPPEERS has been created.

```
R3(config-router)# neighbor IBGPPEERS peer-group
```

Next, assign policies to the peer group.

```
R3(config-router)# neighbor IBGPPEERS remote-as 100
R3(config-router)# neighbor IBGPPEERS route-map INTERNAL out
R3(config-router)# neighbor IBGPPEERS filter-list 1 out
```

```
R3(config-router)# neighbor IBGPPEERS filter-list 2 in
R3(config-router)# neighbor IBGPPEERS next-hop-self
R3(config-router)# neighbor IBGPPEERS soft-reconfiguration in
R3(config-router)# neighbor IBGPPEERS update-source loopback 0
```

Then assign each neighbor to be a member of the peer-group.

```
R3(config-router)# neighbor 4.4.4.4 peer-group IBGPPEERS
R3(config-router)# neighbor 5.5.5.5 peer-group IBGPPEERS
R3(config-router)# neighbor 6.6.6.6 peer-group IBGPPEERS
R3(config-router)# neighbor 1.1.1.1 peer-group IBGPPEERS
```

Finally, R1 had an additional inbound filter list it needed applied.

```
R3(config-router)# neighbor 1.1.1.1 filter-list 3 in
```

Summary

This section summarizes the key points discussed in this lesson.

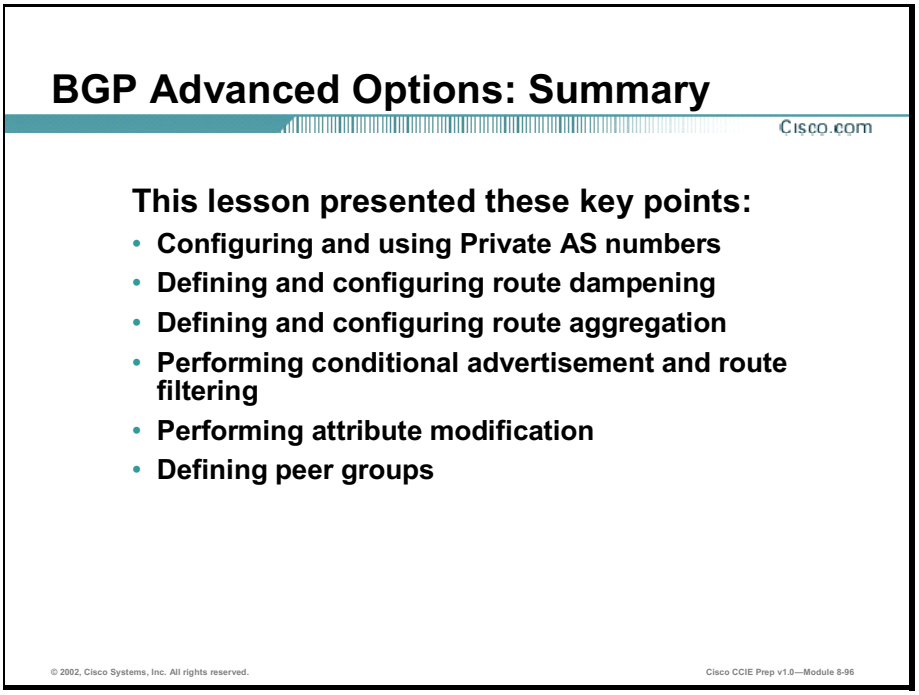

Next Steps

After completing this lesson, go to:

- Troubleshooting

References

For additional information, refer to these resources:

- *Internet Routing Architectures* by Sam Halabi

Lesson Assessment

Q1) Private Autonomous System (AS) numbers fall into which range?

A) 1-1023

B) 1024-2048

C) 65550-65535

D) 64152 to 65535

Q2) What is the proper term that describes when a Border Gateway Protocol (BGP) prefix is constantly updated and withdrawn from the BGP table?

A) convergence

B) route flapping

C) redistribution

D) dampening

Q3) When you wish to perform filtering via Internet Protocol (IP) addresses, which command(s) could you issue?

A) neighbor <*ip-address*> prefix-list

B) neighbor <*ip-address*> distribute-list

C) neighbor <*ip-address*> as-path-list

D) neighbor <*ip-address*> filter-list

Q4) When you wish to perform filtering via an AS path, which command(s) could you issue?

A) neighbor <*ip-address*> prefix-list

B) neighbor <*ip-address*> distribute-list

C) neighbor <*ip-address*> as-path-list

D) neighbor <*ip-address*> filter-list

Q5) Which of the following regular expressions will only allow networks originating from AS 600 to enter a BGP router?

A) ip as-path access-list 1 permit ^600$

B) ip as-path access-list 1 permit $600_

C) ip as-path access-list 1 permit ^600_

D) ip as-path access-list 1 permit _600_

Troubleshooting

Overview

The Border Gateway Protocol (BGP) is a distance-vector routing protocol that is used to exchange routing information among different Autonomous Systems (AS). This lesson will discuss the most common troubleshooting commands to use in a BGP environment.

Importance

BGP version 4 is the current exterior routing protocol used on the Internet. When connecting a multi-homed enterprise to the Internet, it is vital that the administrator know BGP properties, configurations and troubleshooting tips.

Objectives

Upon completing this lesson, you will be able to:

- Issue the proper show command to view relevant information for troubleshooting
- Issue the proper debug command to obtain relevant information for troubleshooting

Learner Skills and Knowledge

To fully benefit from this lesson, you must have these prerequisite skills and knowledge:

- Passed the Cisco Certified Internetworking Expert (CCIE) written qualification exam
- Knowledge of Transfer Control Protocol/Internet Protocol (TCP/IP), access control lists, and Interior Gateway Protocols (IGPs)

Outline

This lesson includes these sections:

- Overview
- Show Commands
- Debug Commands
- Summary
- Lesson Assessment (Quiz)

Show Commands

This section will discuss the common show commands used to view details when troubleshooting BGP.

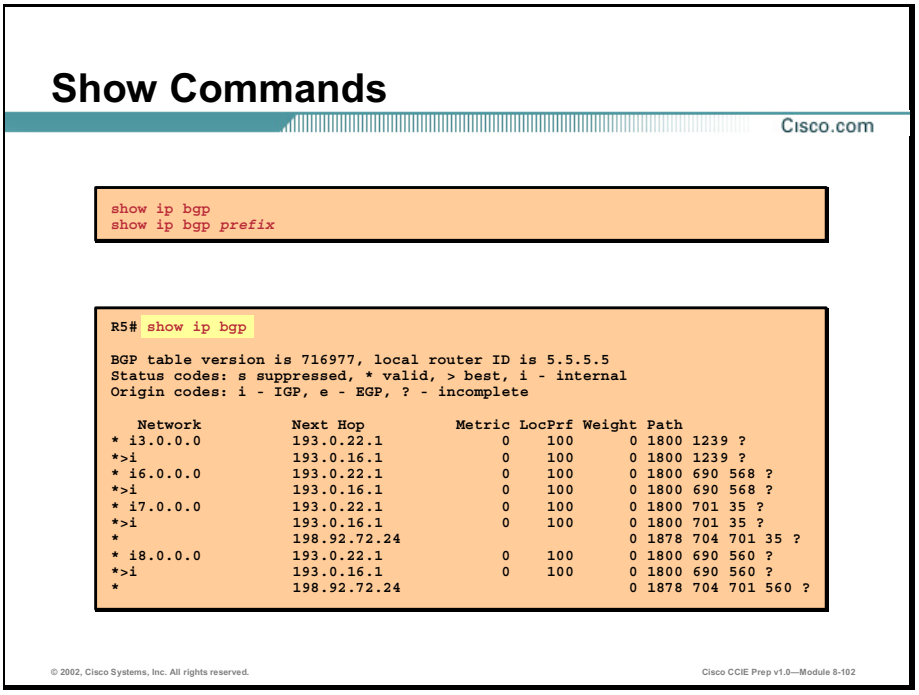

Use the **show ip bgp** command to display entries in the Border Gateway Protocol (BGP) routing table.

```
show ip bgp
show ip bgp prefix
```

The following is sample output from the **show ip bgp** command:

```
R5# show ip bgp

BGP table version is 716977, local router ID is 5.5.5.5
Status codes: s suppressed, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
* i3.0.0.0          193.0.22.1             0    100      0 1800 1239 ?
*>i                 193.0.16.1             0    100      0 1800 1239 ?
* i6.0.0.0          193.0.22.1             0    100      0 1800 690 568 ?
*>i                 193.0.16.1             0    100      0 1800 690 568 ?
* i7.0.0.0          193.0.22.1             0    100      0 1800 701 35 ?
*>i                 193.0.16.1             0    100      0 1800 701 35 ?
```

```
*                 198.92.72.24                          0 1878 704 701 35 ?
* i8.0.0.0        193.0.22.1            0    100      0 1800 690 560 ?
*>i               193.0.16.1            0    100      0 1800 690 560 ?
*                 198.92.72.24                          0 1878 704 701 560 ?
* i13.0.0.0       193.0.22.1            0    100      0 1800 690 200 ?
*>i               193.0.16.1            0    100      0 1800 690 200 ?
*                 198.92.72.24                          0 1878 704 701 200 ?
* i15.0.0.0       193.0.22.1            0    100      0 1800 174 ?
*>i               193.0.16.1            0    100      0 1800 174 ?
* i16.0.0.0       193.0.22.1            0    100      0 1800 701 i
*>i               193.0.16.1            0    100      0 1800 701 i
*                 198.92.72.24                          0 1878 704 701 i
```

The following table describes significant fields shown in the display.

Table 7-3: IP BGP Field Descriptions

Field	Description
`BGP table version`	Internal version number of the table. This number is incremented whenever the table changes
`local router ID`	Internet Protocol (IP) address of the router
`Status codes`	Status of the table entry. The status is displayed at the beginning of each line in the table. It can be one of the following values: s—The table entry is suppressed *—The table entry is valid >—The table entry is the best entry to use for that network i—The table entry was learned via an internal BGP session
`Origin codes`	Indicates the origin of the entry. The origin code is placed at the end of each line in the table. It can be one of the following values: i—Entry originated from Interior Gateway Protocol (IGP) and was advertised with a **network** router configuration command e—Entry originated from Exterior Gateway Protocol (EGP) ?—Origin of the path is not clear. Usually, this is a router that is redistributed into BGP from an IGP
`Network`	IP address of a network entity
`Next Hop`	IP address of the next system that is used when forwarding a packet to the destination network. An entry of 0.0.0.0 indicates that the router has some non-BGP routes to this network.
`Metric`	If shown, this is the value of the interautonomous system metric. This field is frequently not used.
`LocPrf`	Local preference value as set with the **set local-preference** route-map configuration command. The default value is 100.
`Weight`	Weight of the route as set via autonomous system filters
`Path`	Autonomous system paths to the destination network. There can be one entry in this field for each autonomous system in the path.

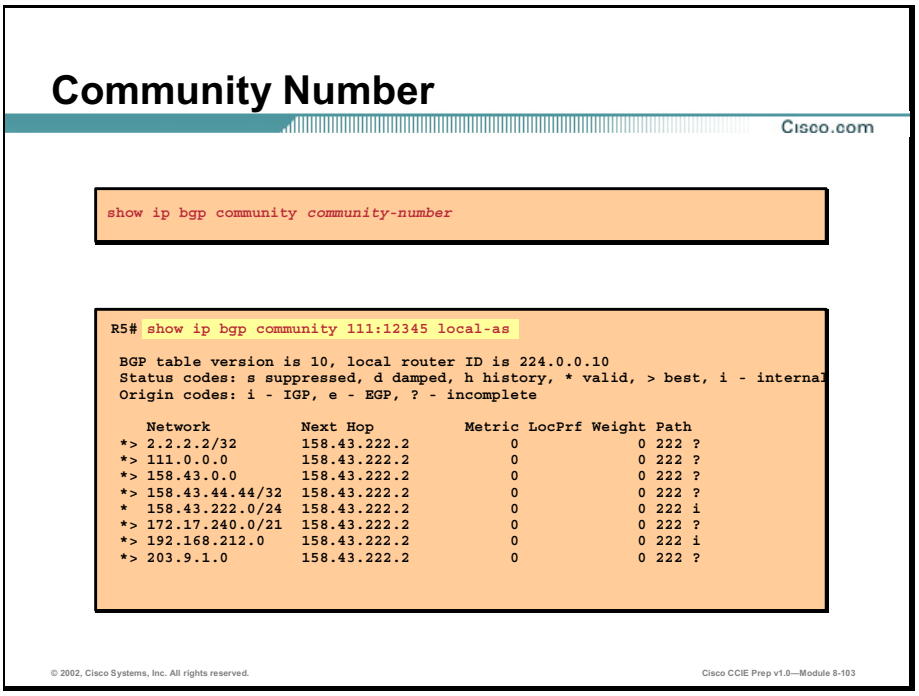

show ip bgp community *community-number(s)*

Use the **show ip bgp community** command to display routes that belong to specified BGP communities. A valid value is a *community-number(s)* in the range 1 to 4294967200, **internet**, **no-export**, **local-as**, or **no-advertise**.

The following is sample output from the **show ip bgp community** command:

```
R5# show ip bgp community 111:12345 local-as

 BGP table version is 10, local router ID is 224.0.0.10
 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
 *> 2.2.2.2/32       158.43.222.2           0             0 222 ?
 *> 111.0.0.0        158.43.222.2           0             0 222 ?
 *> 158.43.0.0       158.43.222.2           0             0 222 ?
 *> 158.43.44.44/32  158.43.222.2           0             0 222 ?
 *  158.43.222.0/24  158.43.222.2           0             0 222 i
 *> 172.17.240.0/21  158.43.222.2           0             0 222 ?
 *> 192.168.212.0    158.43.222.2           0             0 222 i
 *> 203.9.1.0        158.43.222.2           0             0 222 ?
```

The following table describes significant fields shown in the display.

Table 7-4: IP BGP Community Field Descriptions

Field	Description
`BGP table version`	Internal version number of the table. This number is incremented whenever the table changes.
`local router ID`	IP address of the router
`Status codes`	Status of the table entry. The status is displayed at the beginning of each line in the table. It can be one of the following values: s—The table entry is suppressed *—The table entry is valid >—The table entry is the best entry to use for that network i—The table entry was learned via an internal BGP session
`Origin codes`	Indicates the origin of the entry. The origin code is placed at the end of each line in the table. It can be one of the following values: i—Entry originated from IGP and was advertised with a network router configuration command e—Entry originated from EGP ?—Origin of the path is not clear. Usually, this is a router that is redistributed into BGP from an IGP
`Network`	IP address of a network entity
`Next Hop`	IP address of the next system that is used when forwarding a packet to the destination network. An entry of 0.0.0.0 indicates that the router has some non-BGP routes to this network.
`Metric`	If shown, this is the value of the interautonomous system metric. This field is frequently not used.
`LocPrf`	Local preference value as set with the set local-preference route-map configuration command. The default value is 100.
`Weight`	Weight of the route as set via autonomous system filters
`Path`	Autonomous system paths to the destination network. There can be one entry in this field for each autonomous system in the path.

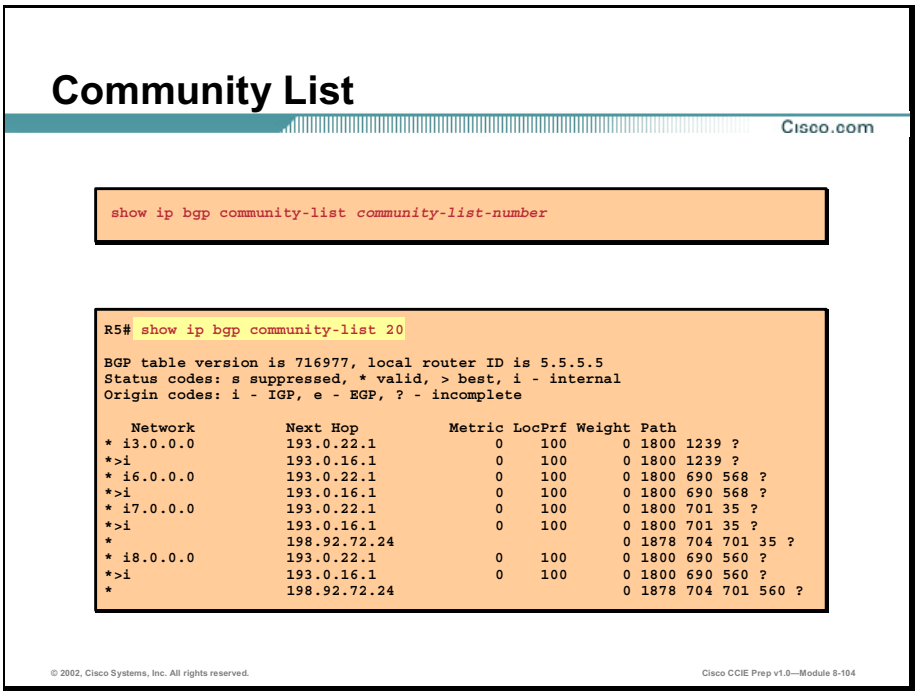

show ip bgp community-list *community-list-number*

Use the **show ip bgp community-list** command to display routes that are permitted by the BGP community list. c*ommunity-list-number* must be a number in the range 1 to 99.

The following is sample output of the **show ip bgp community-list** command:

```
R5# show ip bgp community-list 20

BGP table version is 716977, local router ID is 5.5.5.5
Status codes: s suppressed, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
* i3.0.0.0          193.0.22.1             0    100      0 1800 1239 ?
*>i                 193.0.16.1             0    100      0 1800 1239 ?
* i6.0.0.0          193.0.22.1             0    100      0 1800 690 568 ?
*>i                 193.0.16.1             0    100      0 1800 690 568 ?
* i7.0.0.0          193.0.22.1             0    100      0 1800 701 35 ?
*>i                 193.0.16.1             0    100      0 1800 701 35 ?
*                   198.92.72.24                        0 1878 704 701 35 ?
* i8.0.0.0          193.0.22.1             0    100      0 1800 690 560 ?
*>i                 193.0.16.1             0    100      0 1800 690 560 ?
*                   198.92.72.24                        0 1878 704 701 560 ?
* i13.0.0.0         193.0.22.1             0    100      0 1800 690 200 ?
```

```
*>i                193.0.16.1            0    100      0 1800 690 200 ?
*                  198.92.72.24                        0 1878 704 701 200 ?
* i15.0.0.0        193.0.22.1            0    100      0 1800 174 ?
*>i                193.0.16.1            0    100      0 1800 174 ?
* i16.0.0.0        193.0.22.1            0    100      0 1800 701 i
*>i                193.0.16.1            0    100      0 1800 701 i
*                  198.92.72.24                        0 1878 704 701 i
```

The following table describes significant fields shown in the display.

Table 7-5: IP BGP Community-List Field Descriptions

Field	Description
`BGP table version`	Internal version number of the table. This number is incremented whenever the table changes.
`local router ID`	IP address of the router
`Status codes`	Status of the table entry. The status is displayed at the beginning of each line in the table. It can be one of the following values: s—The table entry is suppressed *—The table entry is valid >—The table entry is the best entry to use for that network i—The table entry was learned via an internal BGP session
`Origin codes`	Indicates the origin of the entry. The origin code is placed at the end of each line in the table. It can be one of the following values: i—Entry originated from IGP and was advertised with a **network** router configuration command e—Entry originated from EGP ?—Origin of the path is not clear. Usually, this is a router that is redistributed into BGP from an IGP
`Network`	IP address of a network entity
`Next Hop`	IP address of the next system that is used when forwarding a packet to the destination network. An entry of 0.0.0.0 indicates that the router has some non-BGP routes to this network.
`Metric`	If shown, this is the value of the interautonomous system metric. This field is frequently not used.
`LocPrf`	Local preference value as set with the **set local-preference** route-map configuration command. The default value is 100.
`Weight`	Weight of the route as set via autonomous system filters
`Path`	Autonomous system paths to the destination network. There can be one entry in this field for each autonomous system in the path.

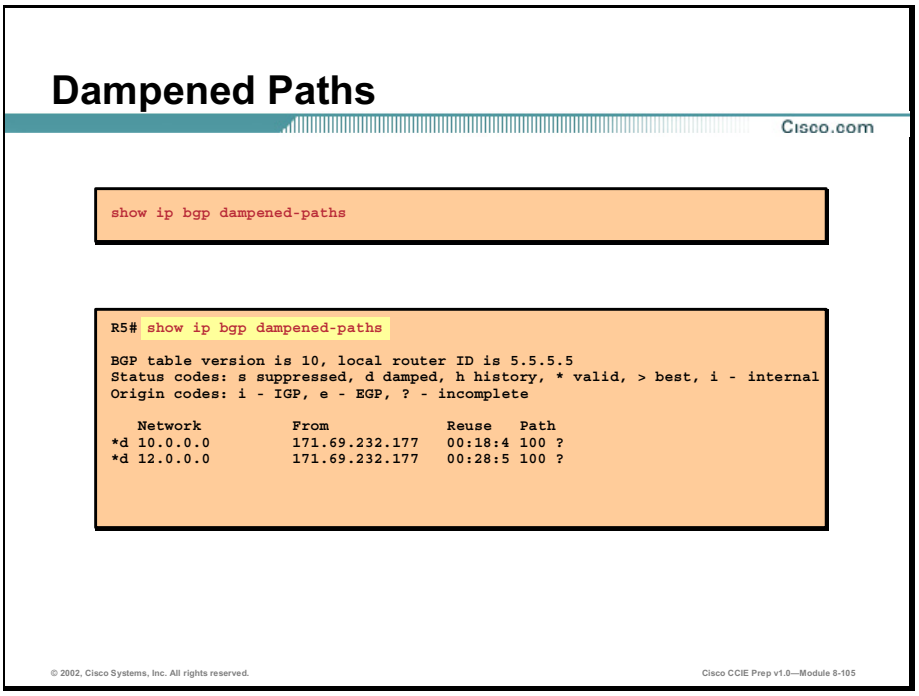

show ip bgp dampened-paths

Use the **show ip bgp dampened-paths** to display BGP dampened routes.

The following is sample output from the **show ip bgp dampened-paths** command:

```
R5# show ip bgp dampened-paths

BGP table version is 10, local router ID is 5.5.5.5
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          From             Reuse    Path
*d 10.0.0.0         171.69.232.177   00:18:4 100 ?
*d 12.0.0.0         171.69.232.177   00:28:5 100 ?
```

The following table describes the fields in the display.

Table 7-6: IP BGP Dampened-Paths Field Descriptions

Field	Description
`BGP table version`	Internal version number for the table. This number is incremented any time the table changes.
`local router`	IP address of the router where route dampening is enabled
`*d Network`	Route to the network indicated is dampened
`From`	IP address of the peer that advertised this path
`Reuse`	Time (in hours:minutes:seconds) after which the path will be made available
`Path`	Autonomous System (AS)-path of the route that is being dampened

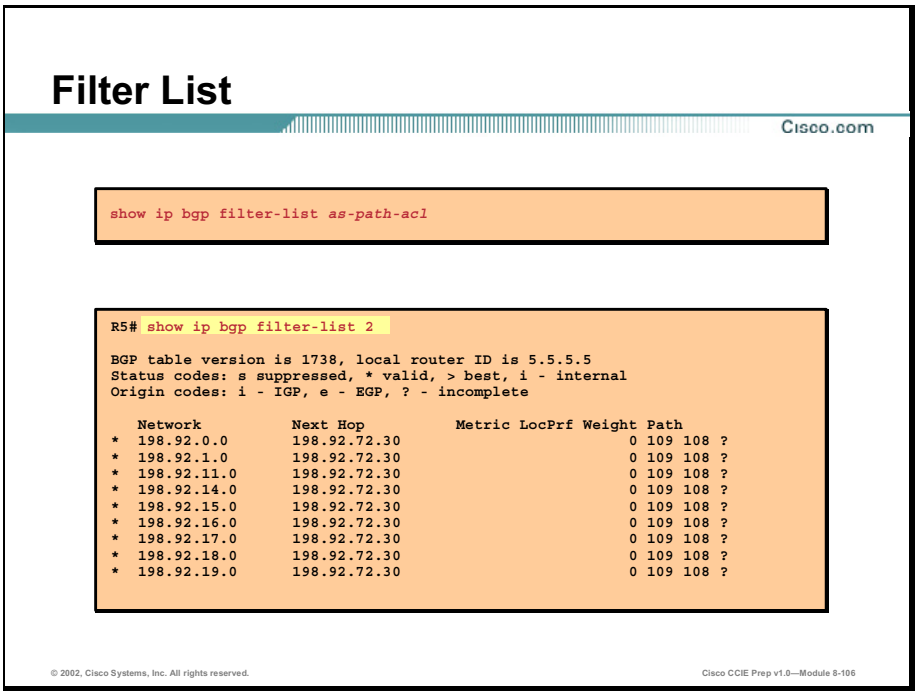

Use the **show ip bgp filter-list** command to display routes that conform to a specified filter list. *as-path-acl* is the number of an autonomous system path access list. It can be a number from 1 to 199.

The following is sample output from the **show ip bgp filter-list** command:

```
R5# show ip bgp filter-list 2

BGP table version is 1738, local router ID is 5.5.5.5
Status codes: s suppressed, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
*  198.92.0.0       198.92.72.30                         0 109 108 ?
*  198.92.1.0       198.92.72.30                         0 109 108 ?
*  198.92.11.0      198.92.72.30                         0 109 108 ?
*  198.92.14.0      198.92.72.30                         0 109 108 ?
*  198.92.15.0      198.92.72.30                         0 109 108 ?
*  198.92.16.0      198.92.72.30                         0 109 108 ?
*  198.92.17.0      198.92.72.30                         0 109 108 ?
*  198.92.18.0      198.92.72.30                         0 109 108 ?
*  198.92.19.0      198.92.72.30                         0 109 108 ?
*  198.92.24.0      198.92.72.30                         0 109 108 ?
*  198.92.29.0      198.92.72.30                         0 109 108 ?
*  198.92.30.0      198.92.72.30                         0 109 108 ?
```

```
*  198.92.33.0      198.92.72.30                          0 109 108 ?
*  198.92.35.0      198.92.72.30                          0 109 108 ?
*  198.92.36.0      198.92.72.30                          0 109 108 ?
*  198.92.37.0      198.92.72.30                          0 109 108 ?
*  198.92.38.0      198.92.72.30                          0 109 108 ?
*  198.92.39.0      198.92.72.30                          0 109 108 ?
```

The following table describes significant fields shown in the display.

Table 7-7: IP BGP Filter-List Field Descriptions

Field	Description
`BGP table version`	Internal version number for the table. This number is incremented any time the table changes.
`local router ID`	An Internet address of the access server
`Status codes`	Indicates the origin of the entry. The origin code is placed at the end of each line in the table. It can be one of the following values: s—The table entry is suppressed *—The table entry is valid >—The table entry is the best entry to use for that network i—The table entry was learned via an internal BGP session
`Origin codes`	Indicates the origin of the entry. The origin code is placed at the end of each line in the table. It can be one of the following values: i—Entry originated from IGP and was advertised with a **network** router configuration command e—Entry originated from EGP ?—Origin of the path is not clear Usually, this is a router that is redistributed into BGP from an IGP
`Network`	Internet address of the network the entry describes.
`Next Hop`	IP address of the next system to use when forwarding a packet to the destination network. An entry of 0.0.0.0 indicates that the access server has some non-BGP route to this network.
`Metric`	If shown, this is the value of the interautonomous system metric. This field is frequently not used.
`LocPrf`	Local preference value. Default is 100.
`Weight`	Set through the use of autonomous system filters
`Path`	Autonomous system paths to the destination network. There can be one entry in this field for each autonomous system in the path. At the end of the path is the origin code for the path. i—The entry was originated with the IGP and advertised with a **network** router configuration command e—The route originated with EGP ?—The origin of the path is not clear. Usually this is a path that is redistributed into BGP from an IGP

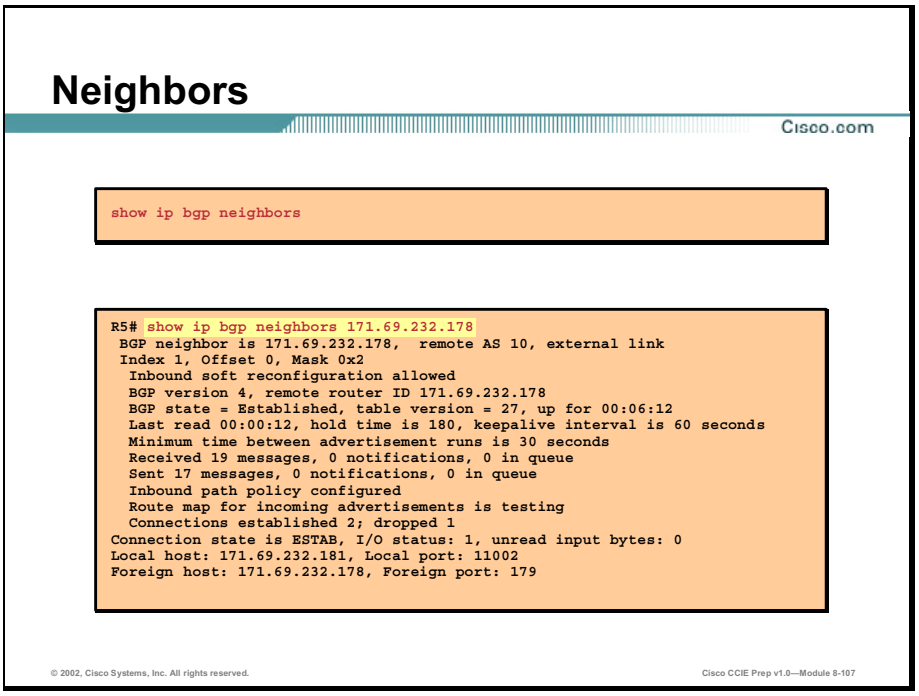

show ip bgp neighbors

Use the **show ip bgp neighbors** command to display information about the TCP and BGP connections to neighbors.

The following is sample output from the **show ip bgp neighbors** command:

```
R5# show ip bgp neighbors 171.69.232.178
 BGP neighbor is 171.69.232.178,  remote AS 10, external link
 Index 1, Offset 0, Mask 0x2
  Inbound soft reconfiguration allowed
  BGP version 4, remote router ID 171.69.232.178
  BGP state = Established, table version = 27, up for 00:06:12
  Last read 00:00:12, hold time is 180, keepalive interval is 60 seconds
  Minimum time between advertisement runs is 30 seconds
  Received 19 messages, 0 notifications, 0 in queue
  Sent 17 messages, 0 notifications, 0 in queue
  Inbound path policy configured
  Route map for incoming advertisements is testing
  Connections established 2; dropped 1
Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Local host: 171.69.232.181, Local port: 11002
Foreign host: 171.69.232.178, Foreign port: 179

Enqueued packets for retransmit: 0, input: 0, saved: 0
```

```
Event Timers (current time is 0x530C294):
Timer          Starts    Wakeups            Next
Retrans            12          0             0x0
TimeWait            0          0             0x0
AckHold            12         10             0x0
SendWnd             0          0             0x0
KeepAlive           0          0             0x0
GiveUp              0          0             0x0
PmtuAger            0          0             0x0

iss:  133981889  snduna:  133982166  sndnxt:  133982166     sndwnd:  16108
irs: 3317025518  rcvnxt: 3317025810  rcvwnd:      16093  delrcvwnd:    291

SRTT: 441 ms, RTTO: 2784 ms, RTV: 951 ms, KRTT: 0 ms
minRTT: 0 ms, maxRTT: 300 ms, ACK hold: 300 ms
Flags: higher precedence, nagle

Datagrams (max data segment is 1460 bytes):
Rcvd: 15 (out of order: 0), with data: 12, total data bytes: 291
Sent: 23 (retransmit: 0), with data: 11, total data bytes: 276
```

The following table describes the fields shown in the display.

Table 7-8: IP BGP Neighbors Field Descriptions

Field	Description
`BGP neighbor`	IP address of the BGP neighbor and its autonomous system number. If the neighbor is in the same autonomous system as the router, then the link between them is internal; otherwise, it is considered external.
`BGP version`	BGP version being used to communicate with the remote router; the neighbor's router ID (an IP address) is also specified
`BGP state`	Internal state of this BGP connection
`table version`	Indicates that the neighbor has been updated with this version of the primary BGP routing table
`up for`	Amount of time that the underlying TCP connection has been in existence
`Last read`	Time that BGP last read a message from this neighbor
`hold time`	Maximum amount of time that can elapse between messages from the peer
`keepalive interval`	Time period between sending keepalive packets, which help ensure that the TCP connection is up
`Received`	Number of total BGP messages received from this peer, including keepalives
`notifications`	Number of error messages received from the peer
`Sent`	Total number of BGP messages that have been sent to this peer, including keepalives

Field	Description
notifications	Number of error messages the router has sent to this peer
Connections established	Number of times the router has established a TCP connection and the two peers have agreed speak BGP with each other
dropped	Number of times that a good connection has failed or been taken down
Connection state	State of BGP peer
unread input bytes	Number of bytes of packets still to be processed
Local host, Local port	Peering address of local router, plus port
Foreign host, Foreign port	Neighbor's peering address
Event Timers	Table displays the number of starts and wakeups for each timer
iss	Initial send sequence number
snduna	Last send sequence number the local host sent but has not received an acknowledgment for
sndnxt	Sequence number the local host will send next
sndwnd	TCP window size of the remote host
irs	Initial receive sequence number
rcvnxt	Last receive sequence number the local host has acknowledged
rcvwnd	Local host's TCP window size
delrecvwnd	Delayed receive window—data the local host has read from the connection, but has not yet subtracted from the receive window the host has advertised to the remote host. The value in this field gradually increases until it is larger than a full-sized packet, at which point it is applied to the rcvwnd field.
SRTT	A calculated smoothed round-trip timeout
RTTO	Round-trip timeout
RTV	Variance of the round-trip time
KRTT	New round-trip timeout (using the Karn algorithm). This field separately tracks the round-trip time of packets that have been retransmitted.
minRTT	Smallest recorded round-trip timeout (hard wire value used for calculation)
maxRTT	Largest recorded round-trip timeout
ACK hold	Time the local host will delay an acknowledgment in order to piggyback data on it
Flags	IP precedence of the BGP packets
Datagrams: Rcvd	Number of update packets received from neighbor
with data	Number of update packets received with data
total data bytes	Total bytes of data
Sent	Number of update packets sent
with data	Number of update packets with data sent
total data bytes	Total number of data bytes

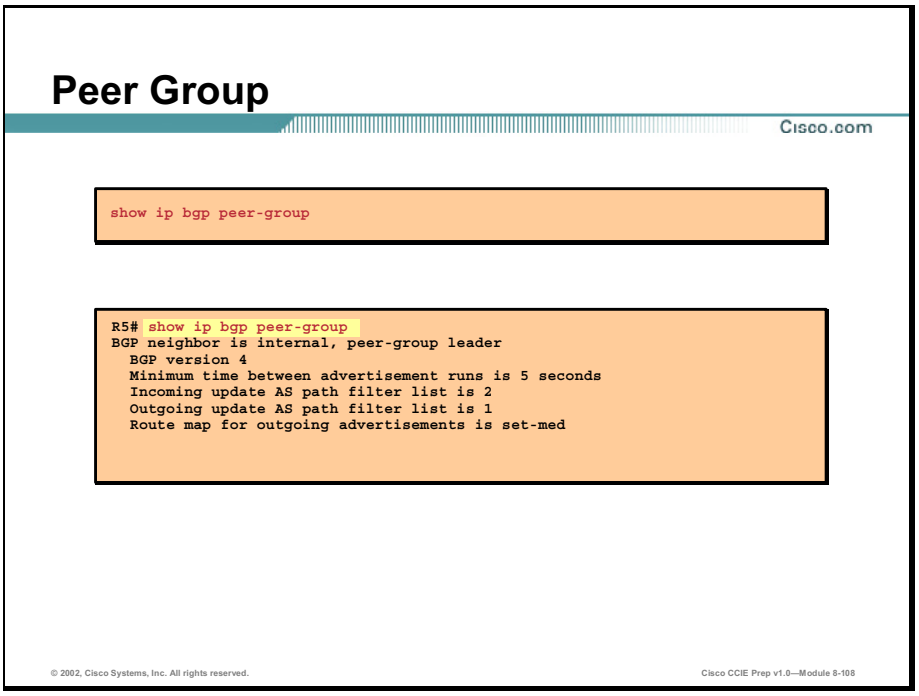

Using Private AS numbers

Use the **show ip bgp peer-group** command to display information about BGP peer groups.

The following is sample output from the **show ip bgp peer-group** command:

```
R5# show ip bgp peer-group

BGP neighbor is internal, peer-group leader
  BGP version 4
  Minimum time between advertisement runs is 5 seconds
  Incoming update AS path filter list is 2
  Outgoing update AS path filter list is 1
  Route map for outgoing advertisements is set-med
```

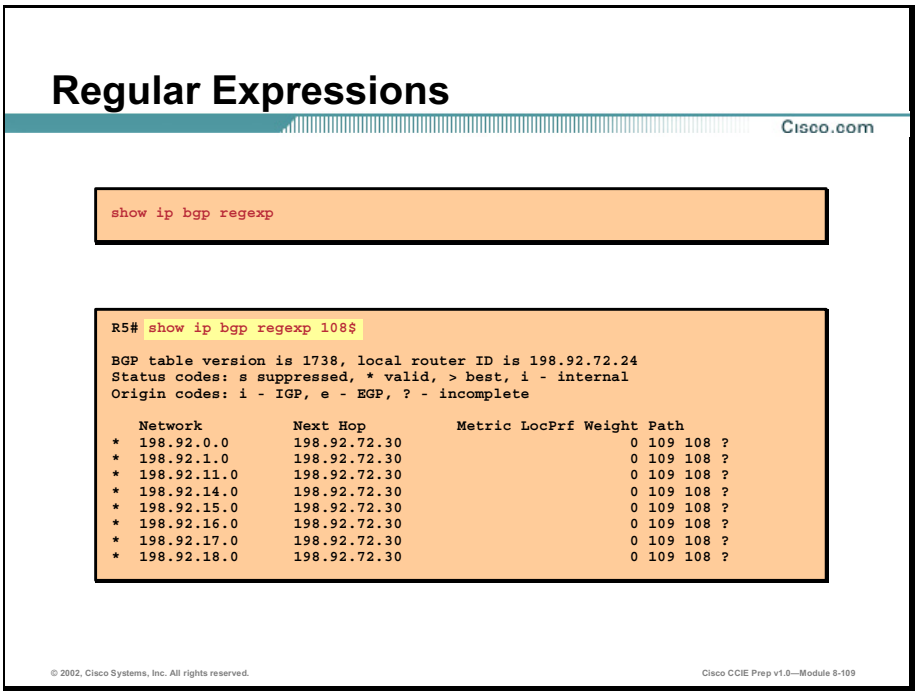

show ip bgp regexp

Use the **show ip bgp regexp** command to display routes matching the regular expression.

The following is sample output from the **show ip bgp regexp** command:

```
R5# show ip bgp regexp 108$

BGP table version is 1738, local router ID is 198.92.72.24
Status codes: s suppressed, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
*  198.92.0.0       198.92.72.30                         0 109 108 ?
*  198.92.1.0       198.92.72.30                         0 109 108 ?
*  198.92.11.0      198.92.72.30                         0 109 108 ?
*  198.92.14.0      198.92.72.30                         0 109 108 ?
*  198.92.15.0      198.92.72.30                         0 109 108 ?
*  198.92.16.0      198.92.72.30                         0 109 108 ?
*  198.92.17.0      198.92.72.30                         0 109 108 ?
*  198.92.18.0      198.92.72.30                         0 109 108 ?
*  198.92.19.0      198.92.72.30                         0 109 108 ?
*  198.92.24.0      198.92.72.30                         0 109 108 ?
*  198.92.29.0      198.92.72.30                         0 109 108 ?
*  198.92.30.0      198.92.72.30                         0 109 108 ?
```

```
*  198.92.33.0      198.92.72.30                    0 109 108 ?
*  198.92.35.0      198.92.72.30                    0 109 108 ?
*  198.92.36.0      198.92.72.30                    0 109 108 ?
*  198.92.37.0      198.92.72.30                    0 109 108 ?
*  198.92.38.0      198.92.72.30                    0 109 108 ?
*  198.92.39.0      198.92.72.30                    0 109 108 ?
```

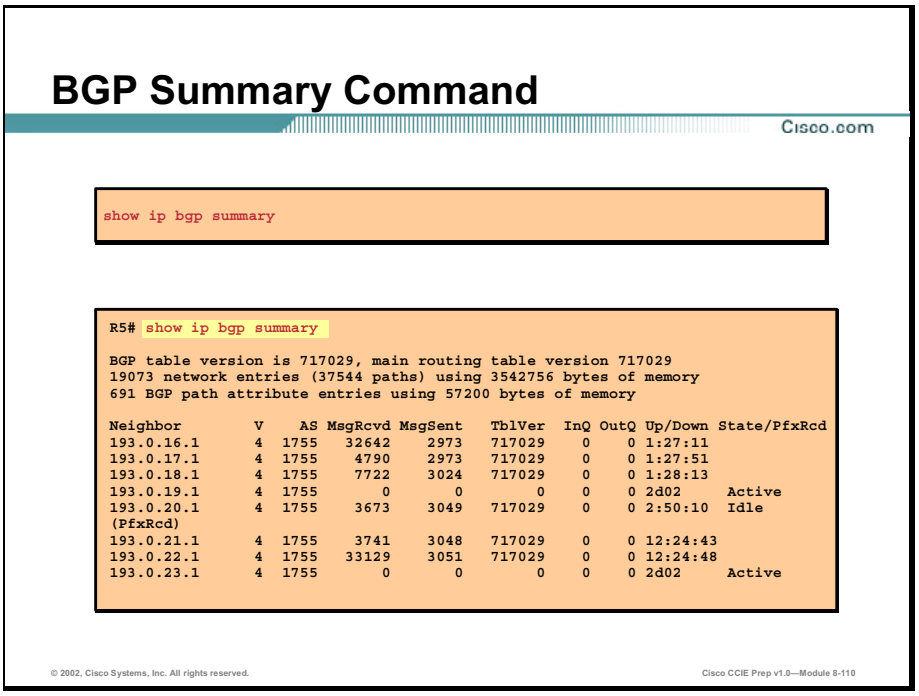

Use the **show ip bgp summary** command to display the status of all BGP connections.

The following is sample output from the **show ip bgp summary** command:

```
R5# show ip bgp summary

BGP table version is 717029, main routing table version 717029
19073 network entries (37544 paths) using 3542756 bytes of memory
691 BGP path attribute entries using 57200 bytes of memory

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
193.0.16.1      4  1755   32642    2973   717029    0    0 1:27:11
193.0.17.1      4  1755    4790    2973   717029    0    0 1:27:51
193.0.18.1      4  1755    7722    3024   717029    0    0 1:28:13
193.0.19.1      4  1755       0       0        0    0    0 2d02     Active
193.0.20.1      4  1755    3673    3049   717029    0    0 2:50:10  Idle
(PfxRcd)
193.0.21.1      4  1755    3741    3048   717029    0    0 12:24:43
193.0.22.1      4  1755   33129    3051   717029    0    0 12:24:48
193.0.23.1      4  1755       0       0        0    0    0 2d02     Active
193.0.24.1      4  1755       0       0        0    0    0 2d02     Active
193.0.25.1      4  1755       0       0        0    0    0 2d02     Active
193.0.26.1      4  1755       0       0        0    0    0 2d02     Active
193.0.27.1      4  1755    4269    3049   717029    0    0 12:39:33
193.0.28.1      4  1755    3037    3050   717029    0    0 2:08:15
```

```
198.92.72.24    4  1878   11635   13300   717028    0    0 0:50:39
198.92.72.36    4  1001       0       0        0    0    0 never    Idle (Admin)
```

The following table describes significant fields shown in the display.

Table 7-9: IP BGP Summary Field Descriptions

Subhead	Subhead
`BGP table version`	Internal version number of BGP database
`main routing table version`	Last version of BGP database that was injected into main routing table
`Neighbor`	IP address of a neighbor
`V`	BGP version number spoken to that neighbor
`AS`	Autonomous System
`MsgRcvd`	BGP messages received from that neighbor
`MsgSent`	BGP messages sent to that neighbor
`TblVer`	Last version of the BGP database that was sent to that neighbor
`InQ`	Number of messages from that neighbor waiting to be processed
`OutQ`	Number of messages waiting to be sent to that neighbor
`Up/Down`	The length of time that the BGP session has been in state Established, or the current state if it is not Established
`State/PfxRcd`	Current state of the BGP session/the number of prefixes the router has received from a neighbor or peer group. When the maximum number (as set by the **neighbor maximum-prefix** command) is reached, the string "PfxRcd" appears in the entry, the neighbor is shut down, and the connection is Idle. An (Admin) entry with Idle status indicates that the connection has been shut down using the **neighbor shutdown** command

Debug Commands

Use the **debug ip bgp** command to display information related to processing BGPs.

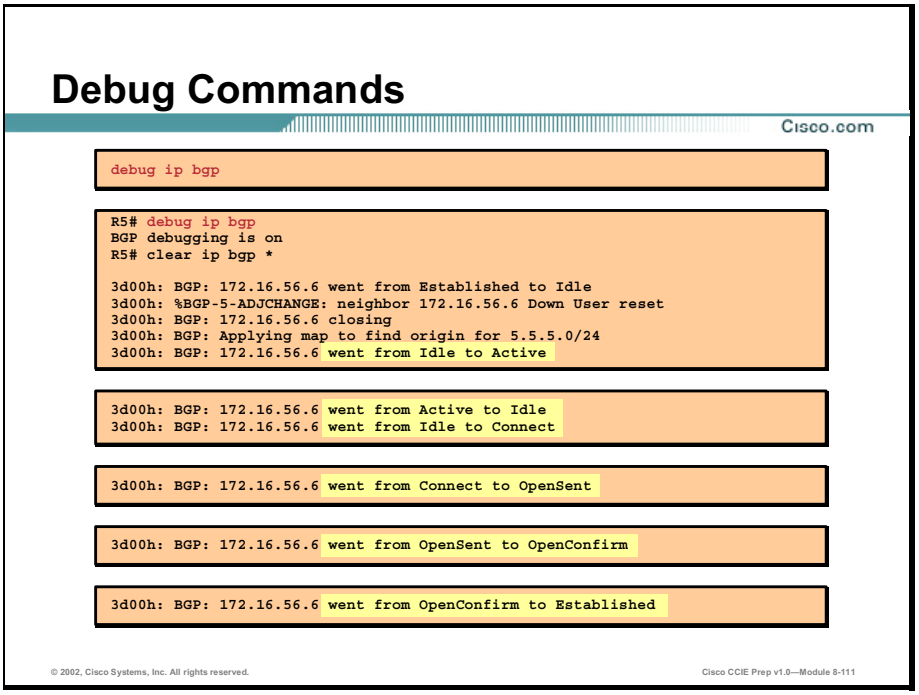

The following is sample output of a BGP speaker making a proper BGP neighbor relationship.

```
R5# debug ip bgp
BGP debugging is on
R5# clear ip bgp *

3d00h: BGP: 172.16.56.6 went from Established to Idle
3d00h: %BGP-5-ADJCHANGE: neighbor 172.16.56.6 Down User reset
3d00h: BGP: 172.16.56.6 closing
3d00h: BGP: Applying map to find origin for 5.5.5.0/24
3d00h: BGP: 172.16.56.6 went from Idle to Active
3d00h: BGP: 172.16.56.6 open active, delay 29472ms
3d00h: BGP: 172.16.56.6 passive open
3d00h: BGP: 172.16.56.6 went from Active to Idle
3d00h: BGP: 172.16.56.6 went from Idle to Connect
3d00h: BGP: 172.16.56.6 rcv message type 1, length (excl. header) 26
3d00h: BGP: 172.16.56.6 rcv OPEN, version 4
3d00h: BGP: 172.16.56.6 went from Connect to OpenSent
3d00h: BGP: 172.16.56.6 sending OPEN, version 4, my as: 500
3d00h: BGP: 172.16.56.6 rcv OPEN w/ OPTION parameter len: 16
```

```
3d00h: BGP: 172.16.56.6 rcvd OPEN w/ optional parameter type 2 (Capability) len
6
3d00h: BGP: 172.16.56.6 OPEN has CAPABILITY code: 1, length 4
3d00h: BGP: 172.16.56.6 OPEN has MP_EXT CAP for afi/safi: 1/1
3d00h: BGP: 172.16.56.6 rcvd OPEN w/ optional parameter type 2 (Capability) len
2
3d00h: BGP: 172.16.56.6 OPEN has CAPABILITY code: 128, length 0
3d00h: BGP: 172.16.56.6 OPEN has ROUTE-REFRESH capability(old) for all address-
families
3d00h: BGP: 172.16.56.6 rcvd OPEN w/ optional parameter type 2 (Capability) len
2
3d00h: BGP: 172.16.56.6 OPEN has CAPABILITY code: 2, length 0
3d00h: BGP: 172.16.56.6 OPEN has ROUTE-REFRESH capability(new) for all address-
families
3d00h: BGP: 172.16.56.6 went from OpenSent to OpenConfirm
3d00h: BGP: 172.16.56.6 send message type 1, length (incl. header) 45
3d00h: BGP: 172.16.56.6 send message type 4, length (incl. header) 19
3d00h: BGP: 172.16.56.6 rcv message type 4, length (excl. header) 0
3d00h: BGP: 172.16.56.6 went from OpenConfirm to Established
3d00h: %BGP-5-ADJCHANGE: neighbor 172.16.56.6 Up
3d00h: BGP: 172.16.56.6 send message type 4, length (incl. header) 19
3d00h: BGP: 172.16.56.6 rcv message type 4, length (excl. header) 0
```

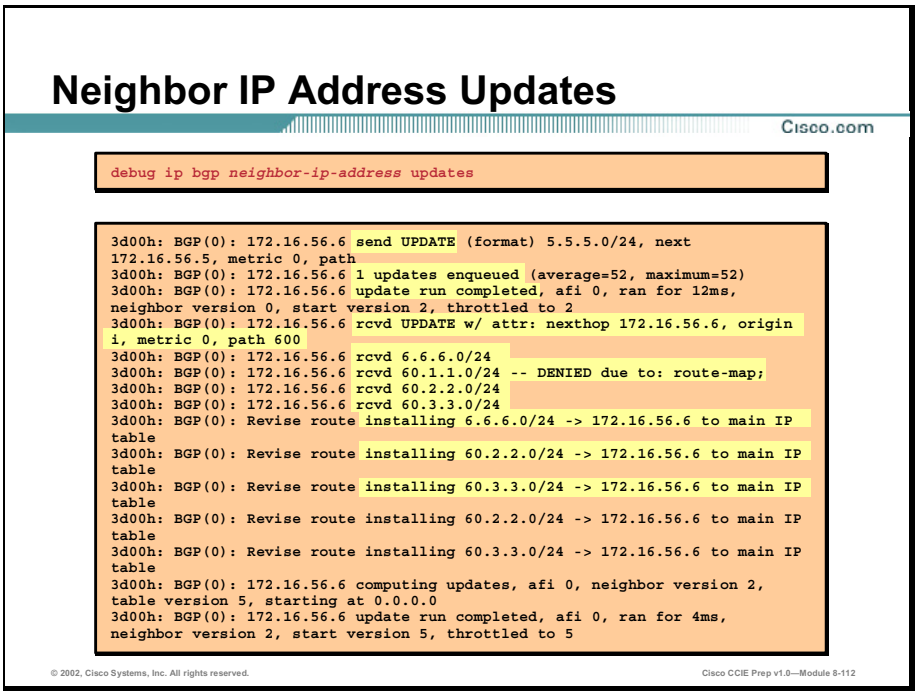

debug ip bgp *neighbor-ip-address* **updates**

Use the **debug ip bgp updates** command to displays BGP updates.

The following is sample output of the **debug ip bgp updates** command.

```
R5# debug ip bgp updates
BGP updates debugging is on
R5# clear ip bgp *
3d00h: %BGP-5-ADJCHANGE: neighbor 172.16.56.6 Down User reset
3d00h: BGP(0): nettable_walker 5.5.5.0/24 route sourced locally
3d00h: %BGP-5-ADJCHANGE: neighbor 172.16.56.6 Up
3d00h: BGP(0): 172.16.56.6 computing updates, afi 0, neighbor version 0, table
version 2, starting at 0.0.0.0
3d00h: BGP(0): 172.16.56.6 send UPDATE (format) 5.5.5.0/24, next 172.16.56.5,
metric 0, path
3d00h: BGP(0): 172.16.56.6 1 updates enqueued (average=52, maximum=52)
3d00h: BGP(0): 172.16.56.6 update run completed, afi 0, ran for 12ms, neighbor
version 0, start version 2, throttled to 2
3d00h: BGP(0): 172.16.56.6 rcvd UPDATE w/ attr: nexthop 172.16.56.6, origin i,
metric 0, path 600
3d00h: BGP(0): 172.16.56.6 rcvd 6.6.6.0/24
3d00h: BGP(0): 172.16.56.6 rcvd 60.1.1.0/24 -- DENIED due to: route-map;
3d00h: BGP(0): 172.16.56.6 rcvd 60.2.2.0/24
3d00h: BGP(0): 172.16.56.6 rcvd 60.3.3.0/24
3d00h: BGP(0): Revise route installing 6.6.6.0/24 -> 172.16.56.6 to main IP
table
```

```
3d00h: BGP(0): Revise route installing 60.2.2.0/24 -> 172.16.56.6 to main IP
table

3d00h: BGP(0): Revise route installing 60.3.3.0/24 -> 172.16.56.6 to main IP
table

3d00h: BGP(0): 172.16.56.6 computing updates, afi 0, neighbor version 2, table
version 5, starting at 0.0.0.0

3d00h: BGP(0): 172.16.56.6 update run completed, afi 0, ran for 4ms, neighbor
version 2, start version 5, throttled to 5
```

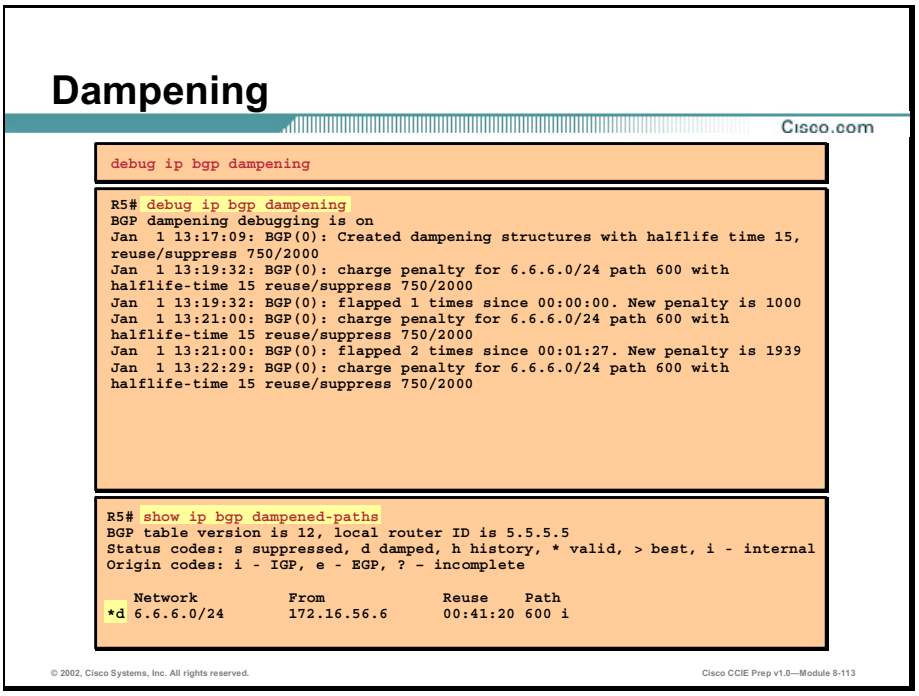

debug ip bgp dampening

Use the **debug ip bgp dampening** command to displays BGP dampening.

The following is sample output of the **debug ip bgp dampening** command.

```
R5# debug ip bgp dampening
BGP dampening debugging is on
Jan  1 13:17:09: BGP(0): Created dampening structures with halflife time 15, reu
se/suppress 750/2000
Jan  1 13:19:32: BGP(0): charge penalty for 6.6.6.0/24 path 600 with halflife-ti
me 15 reuse/suppress 750/2000
Jan  1 13:19:32: BGP(0): flapped 1 times since 00:00:00. New penalty is 1000
Jan  1 13:21:00: BGP(0): charge penalty for 6.6.6.0/24 path 600 with halflife-ti
me 15 reuse/suppress 750/2000
Jan  1 13:21:00: BGP(0): flapped 2 times since 00:01:27. New penalty is 1939
Jan  1 13:22:29: BGP(0): charge penalty for 6.6.6.0/24 path 600 with halflife-ti
me 15 reuse/suppress 750/2000
Jan  1 13:22:29: BGP(0): flapped 3 times since 00:02:56. New penalty is 2821
Jan  1 13:23:12: BGP(0): suppress 6.6.6.0/24 path 600 for 00:28:00 (penalty
2735)
Jan  1 13:23:12: halflife-time 15, reuse/suppress 750/2000
Jan  1 13:23:50: BGP(0): charge penalty for 6.6.6.0/24 path 600 with halflife-ti
me 15 reuse/suppress 750/2000
Jan  1 13:23:50: BGP(0): flapped 4 times since 00:04:17. New penalty is 3661
Jan  1 13:24:43: BGP(0): suppress 6.6.6.0/24 path 600 for 00:33:30 (penalty
3535)
```

```
Jan  1 13:24:43: halflife-time 15, reuse/suppress 750/2000
Jan  1 13:25:15: BGP(0): charge penalty for 6.6.6.0/24 path 600 with halflife-ti
me 15 reuse/suppress 750/2000
Jan  1 13:25:15: BGP(0): flapped 5 times since 00:05:42. New penalty is 4453
Jan  1 13:25:49: BGP(0): suppress 6.6.6.0/24 path 600 for 00:38:00 (penalty
4350)
Jan  1 13:25:49: halflife-time 15, reuse/suppress 750/2000
Jan  1 13:26:18: BGP(0): charge penalty for 6.6.6.0/24 path 600 with halflife-ti
me 15 reuse/suppress 750/2000
Jan  1 13:26:18: BGP(0): flapped 6 times since 00:06:46. New penalty is 5266
Jan  1 13:26:47: BGP(0): suppress 6.6.6.0/24 path 600 for 00:41:50 (penalty
5165)
Jan  1 13:26:47: halflife-time 15, reuse/suppress 750/2000
R5# show ip bgp dampened-paths
BGP table version is 12, local router ID is 5.5.5.5
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? – incomplete

   Network          From             Reuse    Path
*d 6.6.6.0/24       172.16.56.6      00:41:20 600 I
```

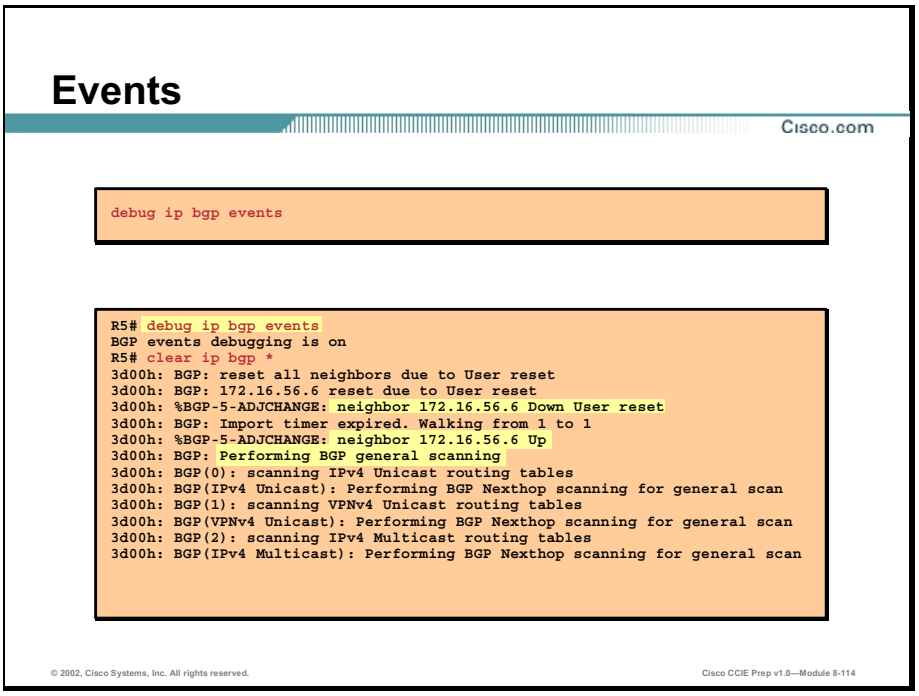

debug ip bgp events

Use the debug ip bgp events commands to display BGP events.

The following is sample output of the **debug ip bgp events** command.

```
R5# debug ip bgp events
BGP events debugging is on
R5# clear ip bgp *
3d00h: BGP: reset all neighbors due to User reset
3d00h: BGP: 172.16.56.6 reset due to User reset
3d00h: %BGP-5-ADJCHANGE: neighbor 172.16.56.6 Down User reset
3d00h: BGP: Import timer expired. Walking from 1 to 1
3d00h: %BGP-5-ADJCHANGE: neighbor 172.16.56.6 Up
3d00h: BGP: Performing BGP general scanning
3d00h: BGP(0): scanning IPv4 Unicast routing tables
3d00h: BGP(IPv4 Unicast): Performing BGP Nexthop scanning for general scan
3d00h: BGP(1): scanning VPNv4 Unicast routing tables
3d00h: BGP(VPNv4 Unicast): Performing BGP Nexthop scanning for general scan
3d00h: BGP(2): scanning IPv4 Multicast routing tables
3d00h: BGP(IPv4 Multicast): Performing BGP Nexthop scanning for general scan
```

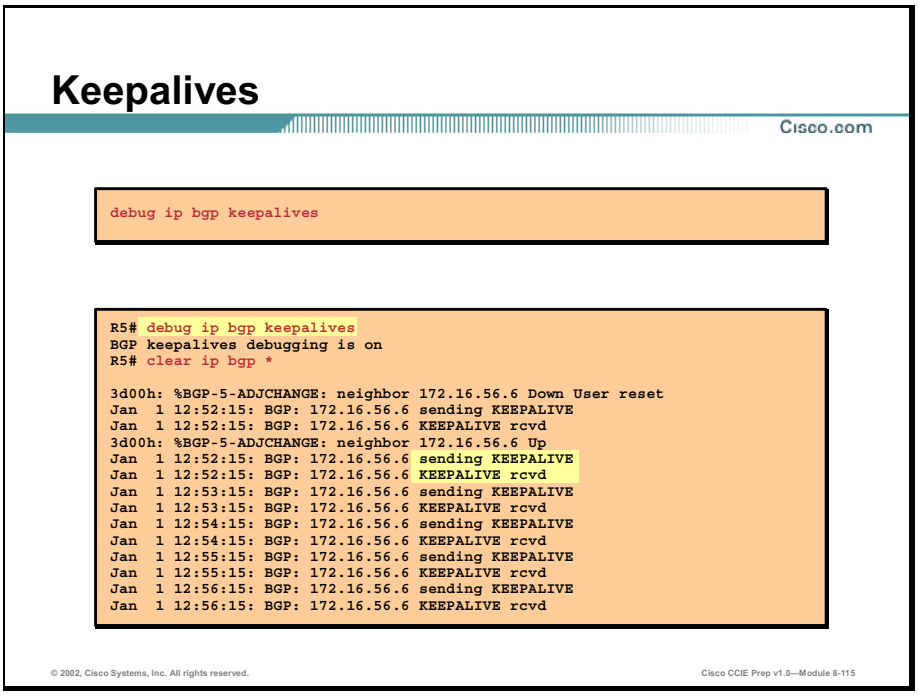

debug ip bgp keepalives

Use the **debug ip bgp keepalives** commands to display BGP keepalives, which by default should occur every 60 seconds.

The following is sample output of the **debug ip bgp keepalives** command.

```
R5# debug ip bgp keepalives
BGP keepalives debugging is on
R5# clear ip bgp *

3d00h: %BGP-5-ADJCHANGE: neighbor 172.16.56.6 Down User reset
Jan  1 12:52:15: BGP: 172.16.56.6 sending KEEPALIVE
Jan  1 12:52:15: BGP: 172.16.56.6 KEEPALIVE rcvd
3d00h: %BGP-5-ADJCHANGE: neighbor 172.16.56.6 Up
Jan  1 12:52:15: BGP: 172.16.56.6 sending KEEPALIVE
Jan  1 12:52:15: BGP: 172.16.56.6 KEEPALIVE rcvd
Jan  1 12:53:15: BGP: 172.16.56.6 sending KEEPALIVE
Jan  1 12:53:15: BGP: 172.16.56.6 KEEPALIVE rcvd
Jan  1 12:54:15: BGP: 172.16.56.6 sending KEEPALIVE
Jan  1 12:54:15: BGP: 172.16.56.6 KEEPALIVE rcvd
Jan  1 12:55:15: BGP: 172.16.56.6 sending KEEPALIVE
Jan  1 12:55:15: BGP: 172.16.56.6 KEEPALIVE rcvd
Jan  1 12:56:15: BGP: 172.16.56.6 sending KEEPALIVE
Jan  1 12:56:15: BGP: 172.16.56.6 KEEPALIVE rcvd
Jan  1 12:57:15: BGP: 172.16.56.6 sending KEEPALIVE
```

```
Jan  1 12:57:15: BGP: 172.16.56.6 KEEPALIVE rcvd
Jan  1 12:58:15: BGP: 172.16.56.6 sending KEEPALIVE
Jan  1 12:58:15: BGP: 172.16.56.6 KEEPALIVE rcvd
```

Summary

This section summarizes the key points discussed in this lesson.

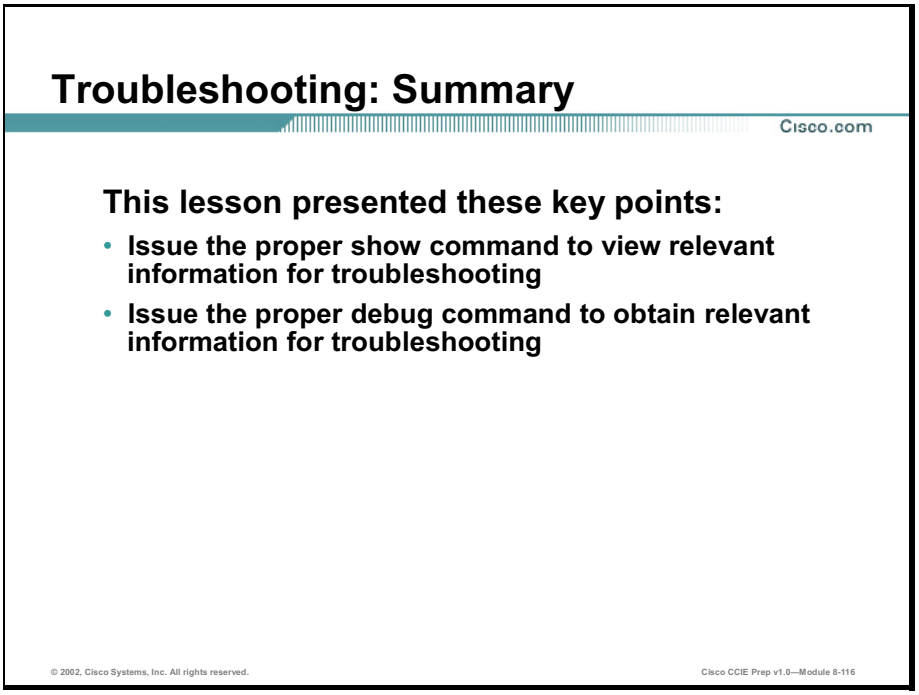

Next Steps

After completing this lesson, go to:

- Advanced Routing Techniques

References

For additional information, refer to these resources:

- *Internet Routing Architectures* by Sam Halabi

Lesson Assessment (Quiz)

Q1) Which command would you issue to display entries in the Border Gateway Protocol (BGP) routing table?

A) show ip route

B) show ip bgp

C) show ip bgp route

D) show ip bgp summary

Q2) Which command would you issue to display routes that belong to specified BGP communities?

A) show ip bgp summary

B) show bgp community

C) show communities

D) show ip bgp community

Q3) Which command would you issue to display information about BGP peer groups?

A) show ip bgp peer group

B) show bgp peer group

C) show ip bgp peer-group

D) show bgp peer-group

Q4) Which debug command would you issue to view output of a BGP speaker making a proper BGP neighbor relationship?

A) show ip bgp

B) debug ip bgp neighbor

C) debug ip bgp

D) debug bgp all

Q5) Which debug command would you issue to display BGP dampening?

A) show ip bgp dampening

B) debug dampening

C) debug ip dampening

D) debug ip bgp dampening

8

Advanced Routing Techniques

Overview

This module covers the redistribution and authentication of multiple routing protocols on a Cisco router.

Upon completing this module, you will be able to:

- Configure static, default, and floating routes
- Configure route redistribution and Policy Routing
- Configure authentication for routing protocols

Outline

The module contains these lessons:

- Static and Default Routing
- Route Redistribution and Control
- Authentication

Static and Default Routing

Overview

Effectively configuring Routing Information Protocol (RIP), Open Shortest Path First (OSPF), Interior Gateway Routing Protocol (IGRP), and Enhanced IGRP (EIGRP) requires a working knowledge of the commands that can set static and default routing. This lesson examines integrating static, floating static, and default routing in a Cisco network.

Importance

Many situations call for static, default, and/or floating static routes for fault tolerance. Understanding the interaction of static and default routes, and their interaction with dynamic routing protocols, such as EIGRP, RIP, and OSPF, is critical to the success of your network.

Objectives

Upon completing this lesson, you will be able to describe:

- The concepts behind a static and floating static route
- The purpose of the default-network and default-information originate commands
- The interaction of the 0.0.0.0 route as it relates to dynamic routing protocols

Learner Skills and Knowledge

To fully benefit from this lesson, you must have these prerequisite skills and knowledge:

- Passed the Cisco Certified Internetworking Expert (CCIE) written qualification exam
- Understand the basic concepts and configuration for RIP, OSPF, IGRP, and EIGRP
- Understand the concepts behind a static route

Outline

This lesson includes these sections:

- Overview
- Static and Floating Routes
- Default Routing
- The Route 0.0.0.0
- Summary
- Lesson Assessment (Quiz)

Static and Floating Routes

This section reviews the syntax for adding a static route, as well as the concept of a “floating” static route.

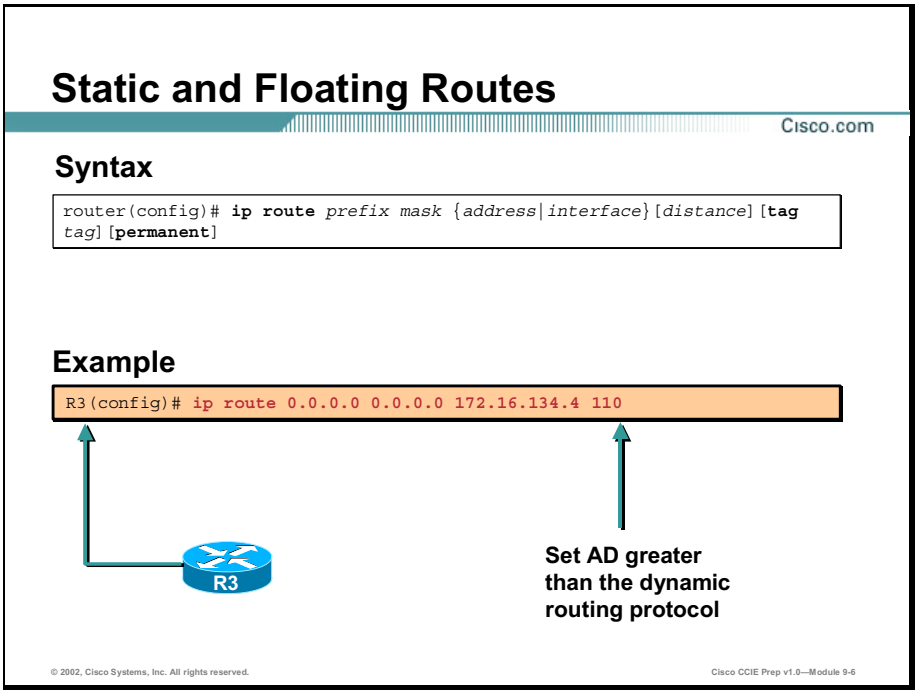

A static route is appropriate when the Cisco Internetwork Operating System (IOS) software cannot dynamically build a route to the destination.

To establish static routes, use the **ip route** global configuration command. To remove static routes, use the **no** form of this command.

ip route *prefix mask* {*address* | *interface*} [*distance*] [**tag** *tag*] [**permanent**]

Table 8-1: <ip route> Commands

Command	Description
distance	(Optional) An administrative distance
tag *tag*	(Optional) Tag value that can be used as a "match" value for controlling redistribution via route maps
permanent	(Optional) Specifies that the route will not be removed, even if the interface shuts down

If you specify an administrative distance, you are flagging a static route that can be overridden by dynamic information. For example, Interior Gateway Routing Protocol (IGRP)-derived routes have a default administrative distance of 100. To have a static route that would be overridden by an IGRP dynamic route, specify an administrative distance greater than 100. This route would then be an example of a floating static route. Administrative distances can be

configured on a static route so that the static route is less desirable than a dynamic route. In this manner, the static route is not used when the dynamic route is available. However, if the dynamic route is lost, the static route can take over, and traffic can be sent through this alternate route. If this alternate route is provided using a Dial-on-Demand Routing (DDR) interface, then that interface can be used as a backup mechanism.

Static routes have a default administrative distance of 1 when the static route points to an Internet Protocol (IP) address, or 0 when it points to an interface.

Static routes that point to an interface will be advertised via Routing Information Protocol (RIP), IGRP, and other dynamic routing protocols, regardless of whether **redistribute static** commands were specified for those routing protocols. This is because static routes that point to an interface are considered to be connected in the routing table and therefore, lose their static nature.

The following example chooses an administrative distance of 110. In this case, packets for network 10.0.0.0 will be routed through to a router at 172.31.3.4 if dynamic information with administrative distance less than 110 is not available.

Example

```
router-3(config)# ip route 10.0.0.0 255.0.0.0 172.31.3.4 110
```

You can create a static route that points to the Null Interface. Enhanced IGRP (EIGRP) always creates a route to a Null interface when it summarizes a group of routes.

Example

```
router(config)# ip route 10.32.0.0 255.255.0.0 Null0 200
```

Default Routing

This section covers how to configure a default route, or gateway of last resort, using the following IP commands: **ip default-gateway** and **ip default-network.**

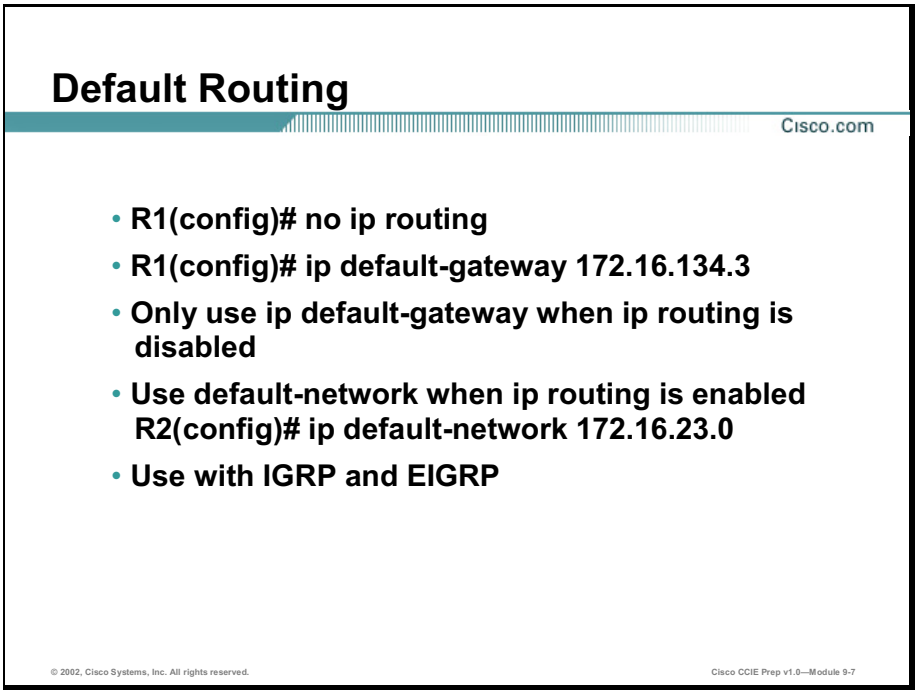

The **ip default-gateway** command differs from the other two commands in that it should only be used when ip routing is disabled on the Cisco router. For instance, if the router is a host in the IP world, you can use this command to define a default gateway for it. You might also use this command when your low end Cisco router is in boot mode in order to Trivial File Transfer Protocol (TFTP) a Cisco IOS® Software image to the router. In boot mode, the router doesn't have ip routing enabled.

Unlike the **ip default-gateway** command, you can use **ip default-network** when ip routing is enabled on the Cisco router. When you configure **ip default-network**, the router considers routes to that network for installation as the gateway of last resort on the router. IP classless must be enabled for a router to forward to a default network. This is enabled with the global configuration command:

```
router(config)# ip classless
```

For every network configured with **ip default-network**, if a router has a route to that network, that route is flagged as a candidate default route. Examine the following routing table taken from a Cisco router:

Example

```
2513# 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
       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, * - candidate default

Gateway of last resort is not set
     161.44.0.0 255.255.255.0 is subnetted, 1 subnets
C       161.44.192.0 is directly connected, Ethernet0
S    198.10.1.0 [1/0] via 161.44.192.2
     131.108.0.0 255.255.255.0 is subnetted, 1 subnets
C       131.108.99.0 is directly connected, TokenRing0
```

Note The static route to 198.10.1.0 is via 161.44.192.2 and the gateway of last resort is not set. If you configure **ip default-network** 198.10.1.0, the routing table changes to the following:

Example

```
2513# 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
       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, * - candidate default

Gateway of last resort is 161.44.192.2 to network 198.10.1.0
     161.44.0.0 255.255.255.0 is subnetted, 1 subnets
C       161.44.192.0 is directly connected, Ethernet0
S       161.44.0.0 255.255.0.0 [1/0] via 161.44.192.0
S*      198.10.1.0 [1/0] via 161.44.192.2
     131.108.0.0 255.255.255.0 is subnetted, 1 subnets
C        131.108.99.0 is directly connected, TokenRing0
```

You can see that the gateway of last resort has now been set as 161.44.192.2. This result is independent of any routing protocol.

Gateways of last resort selected using the **ip default-network** command are propagated differently, depending on which routing protocol is propagating the default route. For IGRP and EIGRP to propagate the route, the network specified by the **ip default-network** command must be known to IGRP or EIGRP. This means the network must be an IGRP- or EIGRP-derived network in the routing table, or the static route used to generate the route to the network must be redistributed into IGRP or EIGRP. For IGRP to advertise a default route, the default route must reside on a different major bit boundary than the interface advertising the route.

The Route 0.0.0.0

This section covers the 0.0.0.0 route and its interaction with the router and its routing protocols.

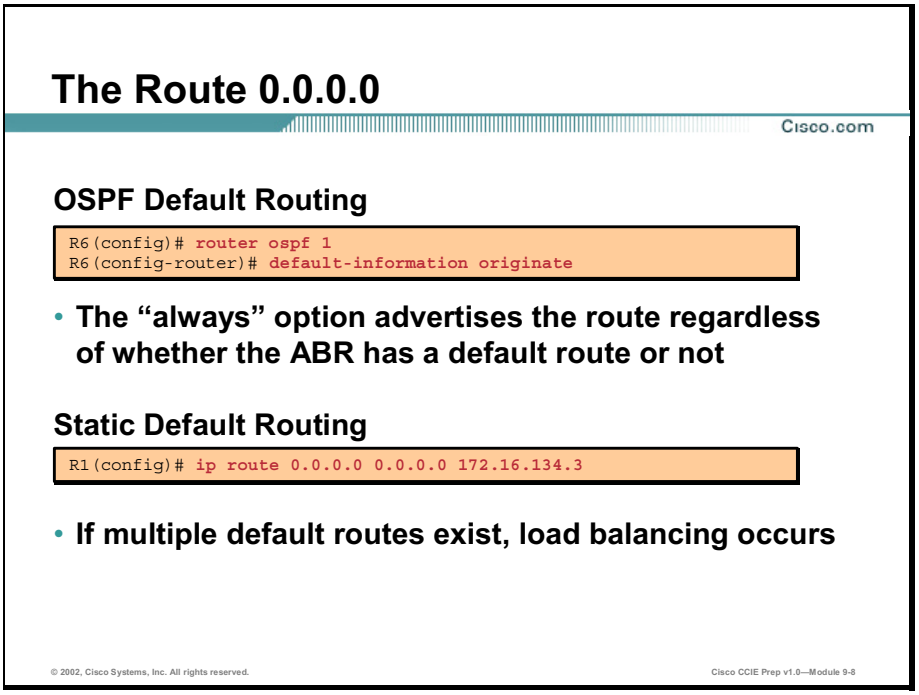

Note	In some IOS releases, RIP does not advertise the default route if the route is not learned via RIP. Therefore, it may be necessary to redistribute the route into RIP, or use the default-information originate command.

RIP advertises a route to 0.0.0.0. OSPF, like RIP, advertises a route for 0.0.0.0 0.0.0.0. However, with OSPF, the router originating the default route must be configured with the **default-information originate** command. The way that OSPF generates default routes (0.0.0.0) varies depending on the type of area the default route is being injected into. This document covers normal areas, stub/totally stubby areas, and Not-So-Stubby Areas (NSSA).

OSPF Normal Areas

By default, normal area routers do not generate default routes. To have an OSPF router generate a default route, use the **default-information originate** [**always**] [**metric** *metric-value*] [**metric-type** *type-value*] [**route-map** *map-name*] command. This generates an external type-2 link with link-state ID 0.0.0.0 and network mask 0.0.0.0, which makes the router an Autonomous System Boundary Router (ASBR).

There are two ways to inject a default route into a normal area. If the ASBR already has the default route, you can advertise 0.0.0.0 into the area. If the ASBR does not have the route, you can add the keyword **always** to the **default-information originate** command, and then advertise 0.0.0.0.

OSPF Stub and Totally Stubby Areas

For stub and totally stubby areas, the Area Border Router (ABR) to the stub area generates a summary Link-State Advertisement (LSA) with the link-state ID 0.0.0.0. This is true even if the ABR does not have a default route. In this scenario, you do not need to use the **default-information originate** command.

OSPF Not-So-Stubby Areas

Creating a static route to network 0.0.0.0 0.0.0.0 is another way to set the gateway of last resort on a router. As with the **ip default-network** command, using the static route to 0.0.0.0 is not dependent on any routing protocols. However, ip routing must be enabled on the router.

Note	IGRP does not understand a route to 0.0.0.0, therefore, it cannot propagate default routes created using the **ip route 0.0.0.0 0.0.0.0** command. Use the **ip default-network** command to have IGRP propagate a default route.

EIGRP. RIP, and OSPF behave as described when using the **ip default-network** command.

Example

Look at an example of configuring a gateway of last resort using the **ip route 0.0.0.0 0.0.0.0** command:

```
router-3# conf terminal
router-3(config)# ip route 0.0.0.0 0.0.0.0 170.170.3.4
router-3(config)# ^Z
router-3#

router-3# 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, * - candidate default
U - per-user static route, o - ODR

Gateway of last resort is 170.170.3.4 to network 0.0.0.0

170.170.0.0/24 is subnetted, 2 subnets
 C 170.170.2.0 is directly connected, Serial0
 C 170.170.3.0 is directly connected, Ethernet0
 S* 0.0.0.0/0 [1/0] via 170.170.3.4
router-3#
```

If you use multiple **ip route 0.0.0.0 0.0.0.0** commands to configure a default route, traffic is load-balanced over the multiple routes.

There are two ways to inject a default route into EIGRP: redistribute a static route or summarize to 0.0.0.0/0. Use the first method when you want to draw all traffic to unknown destinations to a default route at the core of the network. This method is effective for advertising connections to the Internet.

Example

```
ip route 0.0.0.0 0.0.0.0 a.b.c.d (next hop to the internet)
!
router eigrp 100
 redistribute static
 default-metric 10000 10 255 1 1500
```

The static route that is redistributed into EIGRP does not have to be to network 0.0.0.0. If you use another network, you must use the **ip default-network** command to mark the network as a default network. Please refer to Configuring a Gateway of Last Resort for further information.

Summarizing to a default route is effective only when you want to provide remote sites with a default route. Since summaries are configured per interface, you do not need to worry about using distribute-lists or other mechanisms to prevent the default route from being propagated toward the core of your network. Note that a summary to 0.0.0.0/0 overrides a default route learned from any other routing protocol. The only way to configure a default route on a router using this method is to configure a static route to 0.0.0.0/0. (Beginning in Cisco IOS Software 12.0(4)T, you can also configure an administrative distance on the end of the summary-address command, so the local summary doesn't override the 0.0.0.0/0 route).

Example

```
router eigrp 100
 network 10.0.0.0
!
interface serial 0
 encapsulation frame-relay
 no ip address
!
interface serial 0.1 point-to-point
 ip address 10.1.1.1
 frame-relay interface-dlci 10
 ip summary-address eigrp 100 0.0.0.0 0.0.0.0
```

Note	The EIGRP summary-address command also limits or bounds the EIGRP query range, which can be significant in large networks.

Summary

This section summarizes the key points discussed in this lesson.

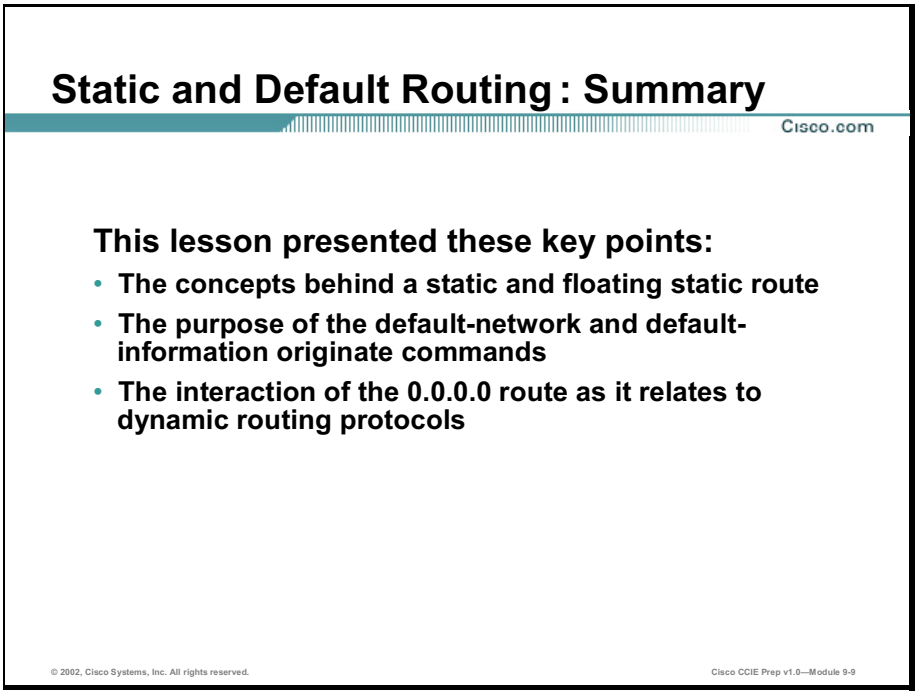

In summary, use the **ip default-gateway** command when ip routing is disabled on a Cisco router. Use the **ip default-network** and **ip route 0.0.0.0 0.0.0.0** commands to set the gateway of last resort on Cisco routers that have ip routing enabled. The way in which routing protocols propagate the default route information varies for each protocol.

Next Steps

After completing this lesson, go to:

- Route Redistribution and Control

References

For additional information, refer to these resources:

- http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fipr_c/ipcprt2/1cfindep.htm

Lesson Assessment (Quiz)

Q1) How can you inject a default route into OSPF?

A) Create a static default route, then redistribute it into OSPF

B) Use the OSPF **default-information originate always** command

C) Create a static default route, and it will automatically find its way into OSPF

D) Create an ABR, and the default route will automatically be injected into the non-backbone area

Q2) Router1 is directly connected to the 135.10.2.0/24 subnet. When router1 pings the address of 135.10.3.1, there is no echo reply. What may cause this problem?

A) No default gateway on source or destination

B) Routing problem somewhere between the two devices

C) Router1 has a default route, and the command **no ip classless** is in the configuration

D) There is no remote device that is running IP with the address of 135.10.3.1

Q3) How can you inject a default route into RIP?

A) Use the RIP **default-information originate** command

B) Create a static default route, and it will automatically find its way into RIP

C) RIP does not support advertisement of the default route

D) Use the **ip default-gateway** command

Route Redistribution and Control

Overview

It is sometimes necessary to accommodate more complex network topologies, such as independent Routing Information Protocol (RIP) and Open Shortest Path First (OSPF) clouds, which must perform mutual redistribution. This lesson will review redistribution between multiple routing protocols, including the concepts of administrative distance and default metrics.

Importance

Understanding how to perform redistribution that allows optimal paths through the network, as well as avoiding routing loops, is critical for the success of the network.

Objectives

Upon completing this lesson, you will be able to:

- Describe route redistribution
- Use the “Default-Metric” within a routing protocol
- Describe the procedure for Variable Length Subnet Mask (VLSM) to Fixed Length Subnet Mask (FLSM) redistribution
- Describe how to filter routes from specific routing protocols

Learner Skills and Knowledge

To fully benefit from this lesson, you must have these prerequisite skills and knowledge:

- Passed the Cisco Certified Internetworking Expert (CCIE) written qualification exam
- Basic understanding of Enhanced Interior Gateway Routing Protocol (EIGRP), Open Shortest Path First (OSPF), Interior Gateway Routing Protocol (IGRP) and Routing Information Protocol (RIP) version 1 and 2
- Understand the concept of administrative distance

Outline

This lesson includes these sections:

- Overview
- Redistribution Review
- Default Metric
- VLSM to FLSM Redistribution
- Summarization
- Filtering
- Summary
- Lesson Assessment (Quiz)

Redistribution Review

This section reviews redistribution concepts.

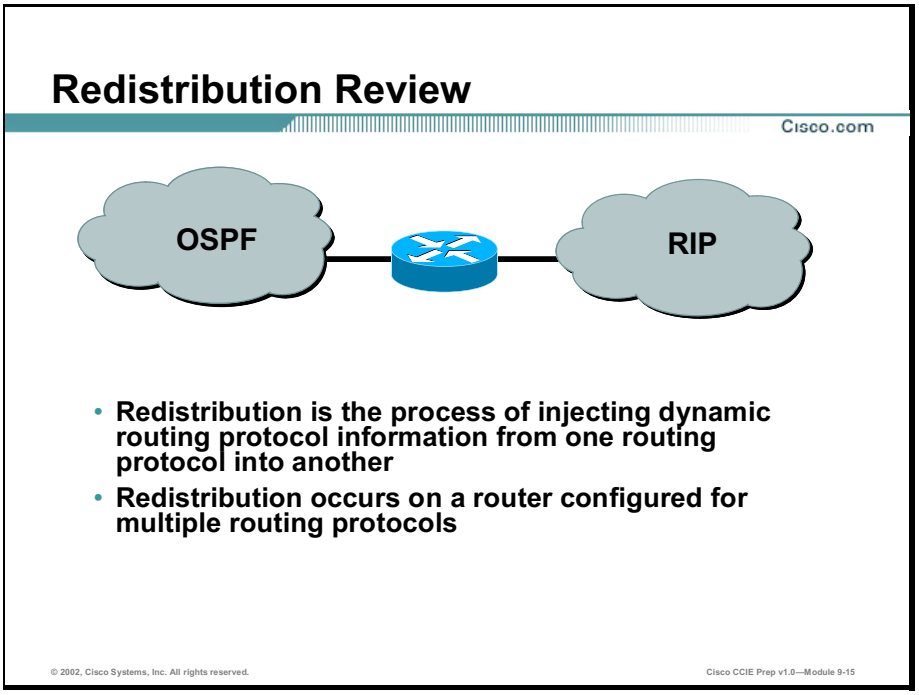

Using a routing protocol to advertise routes that are learned by some other means, such as by another routing protocol, static routes, or directly connected routes, is called redistribution. While running a single routing protocol throughout your entire Internet Protocol (IP) internetwork is desirable, multi-protocol routing is common for a number of reasons, including company mergers, multiple departments managed by multiple network administrators, and multi-vendor environments. Often, running different routing protocols is part of a network design. In any case, having a multiple protocol environment makes redistribution a necessity.

Differences in routing protocol characteristics, such as metrics, administrative distance, and classful and classless capabilities can affect redistribution. Consideration must be given to these differences in order for redistribution to be successful.

If a router is running more than one routing protocol and learns a route to the same destination using both routing protocols, then which route should be selected as the best route? Each protocol uses its own metric type to determine the best route. Comparing routes with different metric types cannot be done. Administrative distances take care of this problem. Administrative distances are assigned to route sources so that the route from the most preferred source will be chosen as the best path.

Default Metric

This section examines the importance of a default metric as it applies to route redistribution.

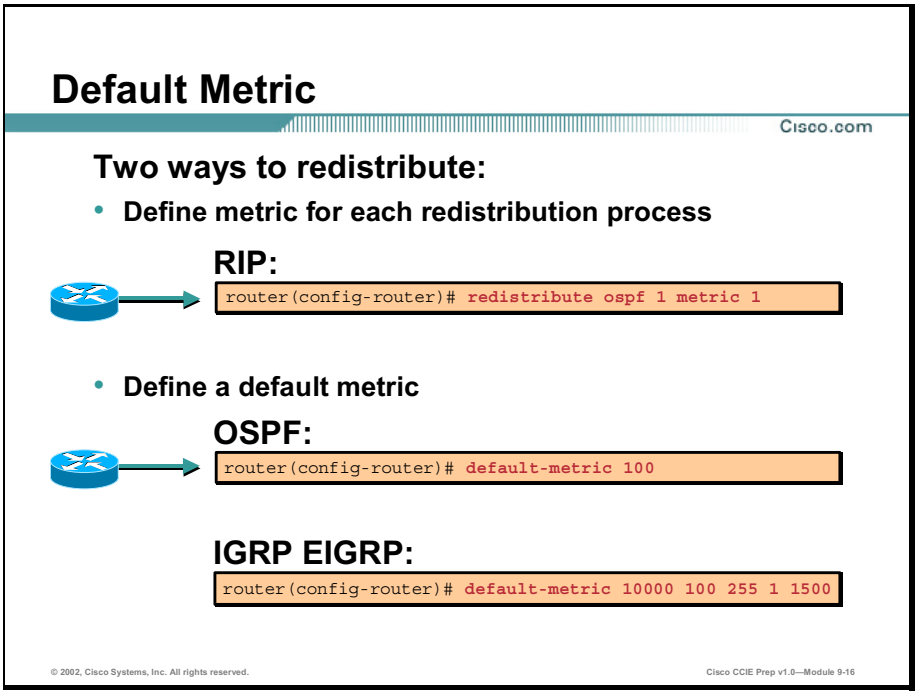

When you redistribute one protocol into another, remember that the metrics of each protocol play an important role in redistribution. Each protocol uses different metrics. For example, the RIP metric is based on hop count while the IGRP/EIGRP metric is based on bandwidth and delay. When routes are redistributed, you must define a metric that is understandable to the receiving protocol. There are two ways of defining metrics when redistributing routes: define the metric for that specific redistribution or use the same metric as a default for all redistribution.

Example

Define the metric for that specific redistribution only:

```
router rip
redistribute static metric 1
redistribute ospf 1 metric 3
```

Example

Alternately, you can use the same metric as a default for all redistribution (Using the **default-metric** command saves work since it eliminates the need for defining the metric separately for each redistribution.)

```
router rip
 redistribute static
 redistribute ospf 1
 default-metric 1
```

IGRP/EIGRP Example

The output shows an IGRP/EIGRP router redistributing static, OSPF, RIP, and IS-IS routes.

```
router igrp/eigrp 1
 network 131.108.0.0
 redistribute static
 redistribute ospf 1
 redistribute rip
 redistribute isis
 default-metric 10000 100 255 1 1500
```

IGRP and EIGRP need five metrics when redistributing other protocols: bandwidth, delay, reliability, load, and Maximum Transmission Unit (MTU) respectively.

The redistribution of IGRP/EIGRP into another IGRP/EIGRP process does not require any metric conversion, so there is no need to define metrics or use the **default-metric** command during redistribution.

OSPF Example

The output shows an OSPF router redistributing static, RIP, IGRP, EIGRP, and IS-IS routes.

```
router ospf 1
 network 131.108.0.0 0.0.255.255 area 0
 redistribute static metric 200 subnets
 redistribute rip metric 200 subnets
 redistribute igrp 1 metric 100 subnets
 redistribute eigrp 1 metric 100 subnets
 redistribute isis metric 10 subnets
```

Note	If no metric is specified, OSPF assigns a default value of 20 when redistributing routes from all protocols except Border Gateway Protocol (BGP) routes, which get a metric of 1.

Whenever there is a major net that is subnetted, you need to use the keyword *subnet* to redistribute protocols into OSPF. Without this keyword, OSPF only redistributes major nets

that are not subnetted. For instance, 131.108.0.0/16 is redistributed without the keyword *subnet*, but 141.108.100.0/24 is not redistributed until you add the *subnet* keyword.

The following output shows a RIP router redistributing static, IGRP, EIGRP, OSPF, and IS-IS routes:

RIP Example

```
router rip
 network 131.108.0.0
 redistribute static
 redistribute igrp 1
 redistribute eigrp 1
 redistribute ospf 1
 redistribute isis
 default-metric 1
```

The RIP metric is composed of hop count and the maximum valid metric is 15. Anything above 15 is considered infinite; you can use 16 to describe an infinite metric in RIP. If we define a metric of 10 for redistributed routes, these routes can only be advertised to routers up to 5 hops away, at which point the metric (hop count) exceeds 15. By defining a metric of 1, you enable a route to travel the maximum number of hops in a RIP domain.

IS-IS Example

```
router isis
network 49.1234.1111.1111.1111.00
redistribute static metric 20
redistribute rip metric 20
redistribute igrp 1 metric 20
redistribute eigrp 1 metric 20
redistribute ospf 1 metric 20
```

The IS-IS metric needs to be between 1 and 63. There is no default-metric option in IS-IS, so you should define a metric for each protocol, as shown in the example.

VLSM to FLSM Redistribution

Classful routing protocols, such as IGRP and RIPv1, do not support variable length subnet masks. This section examines how to redistribute from a classful routing protocol that may have Variable Length Subnet Masks (VLSMs), into a routing protocol that does not support them.

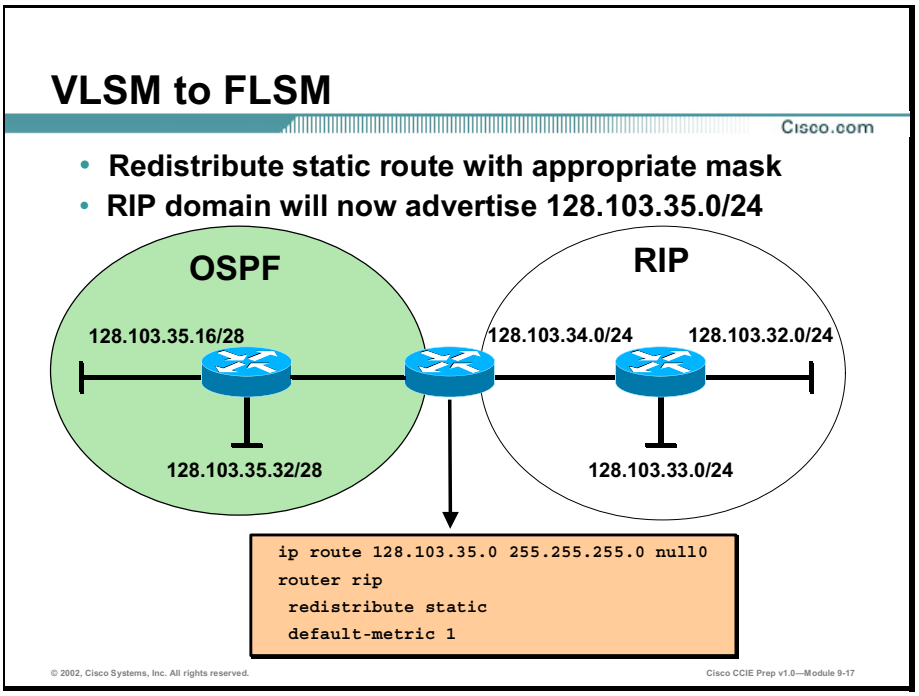

If a router is redistributing between RIP and OSPF and the OSPF domain has a different mask (RIP has 24 bit and OSPF has 28-bit) than the RIP domain, and they are on the same major network, RIP will not redistribute routes learned from OSPF into RIP.

The problem is that RIPv1 (as well as IGRP) do not understand VLSM. The classless routing protocols need to condescend to RIPv1 and IGRP in order for them to redistribute the routes.

One solution is to add a static route in the redistribution router that points to the OSPF domain with a mask of 255.255.255.0, but with a next hop of null0. Then, redistribute static routes into RIP.

Example

```
ip route 128.103.0.0 255.255.0.0 null0
 router rip
 redistribute static
 default metric 1
```

This allows the OSPF route of 128.103.35.0 to be advertised through RIP.

If RIP has a longer mask than OSPF (RIP has 29 bit, OSPF has 28-bit), RIP will not redistribute routes learned from OSPF into RIP. Again, you could add a static route in the redistribution router that points to the OSPF domain with a mask of 255.255.255.248. This way, static routes are redistributed into RIP.

The above solutions also work when you use EIGRP instead of OSPF and IGRP instead of RIP. This problem should not happen if the masks of both protocols are the same or if all the protocols you are using support VLSM. This fix is only considered a patch to cover the RIP and IGRP (VLSM) limitation.

Summarization

This section will cover how to accomplish redistribution of VLSM networks into a classful routing protocol using the technique of summarization.

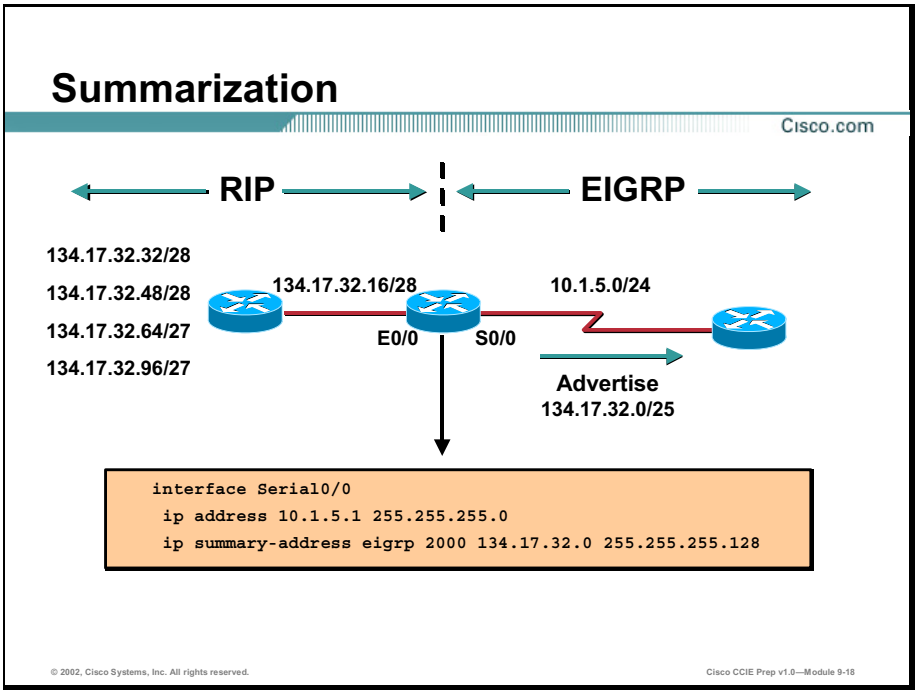

One of the challenges using the static route solutions is that you may be required to redistribute VLSM into classful routing protocols without using static routes. To accomplish the redistribution without using static routes, you need to use summarization.

EIGRP, as well as RIPv2, performs an auto-summarization each time it crosses a border between two different major networks.

EIGRP allows you to summarize internal and external routes on virtually any bit boundary using manual summarization. For example, an EIGRP router could summarize the 192.1.1.0/24, 192.1.2.0/24, and 192.1.3.0/24 into the CIDR block 192.1.0.0/22. You would want to summarize to a classful boundry for the classful routing protocols. By summarizing an EIGRP route to a bit boundry that is acceptable to RIPv1 or IGRP, it will allow the router to advertise the route in the RIPv1 or IGRP domain. In this example, you are summarizing the network of 134.17.32.16/28 down to a 24bit boundry (the one being used by RIP in this example).

Example

```
router# show run
....
!
interface Serial0/0
 ip address 10.1.50.1 255.255.255.0
 ip summary-address eigrp 2000 134.17.32.0 255.255.255.0
!
interface Ethernet0/0
 ip address 134.17.32.17 255.255.255.240
....

router# show ip eigrp topology
IP-EIGRP Topology Table for process 2000

Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - Reply status

P 10.1.50.0/24, 1 successors, FD is 2169856
         via Connected, Serial0/0
P 134.17.32.16/28, 1 successors, FD is 2169856
         via Connected, Ethernet0/0

P 134.17.32.0/24, 1 successors, FD is 10511872
         via Summary (10511872/0), Null0
```

In OSPF, you have 2 ways to summarize the network: using the **area range** option for OSPF routes and the **summary-address** option for redistributed routes.

To specify an address range on an Available Bit Rate (ABR), use the following command in router configuration mode:

area *area-id* **range** *address mask* **[advertise** | **not-advertise]**

Specify an address range for which a single route will be advertised.

To have the OSPF advertise one summary route for all **redistributed** routes covered by a network address and mask, use the following command in router configuration mode:

summary-address *address mask prefix mask* [**not-advertise**] [**tag** *tag*]

Use the optional [**not-advertise**] keyword to filter out a set of routes.

Filtering

This section covers the various methods for controlling what routes are shared among routers using internal routing protocols.

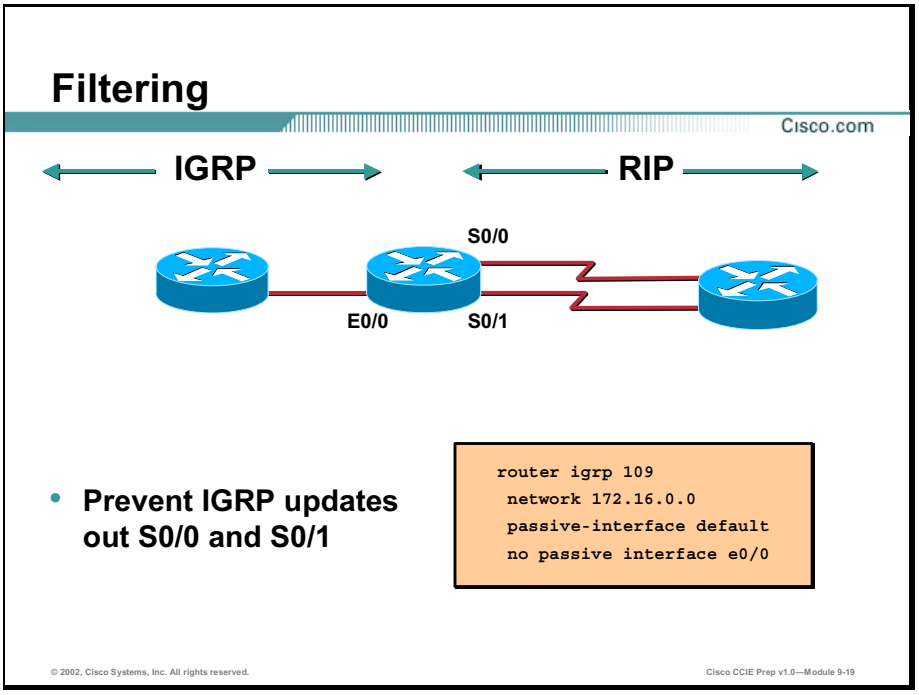

To prevent other routers on a local network from learning about routes dynamically, you can keep routing update messages from being sent through a router interface.

To prevent routing updates through a specified interface, use the following command in router configuration mode:

```
router(config-router)# passive-interface interface-type interface-number
```

You can specify passive for the default on all interfaces.

```
router(config-router)# passive-interface default
```

To activate only those interfaces that need to have adjacencies set, use the following.

```
router(config-router)# no passive-interface
interface-type
```

This will allow the interface to send routing updates.

Example

```
router igrp 109
 network 131.108.0.0
 passive-interface default
 no passive interface e0/0
```

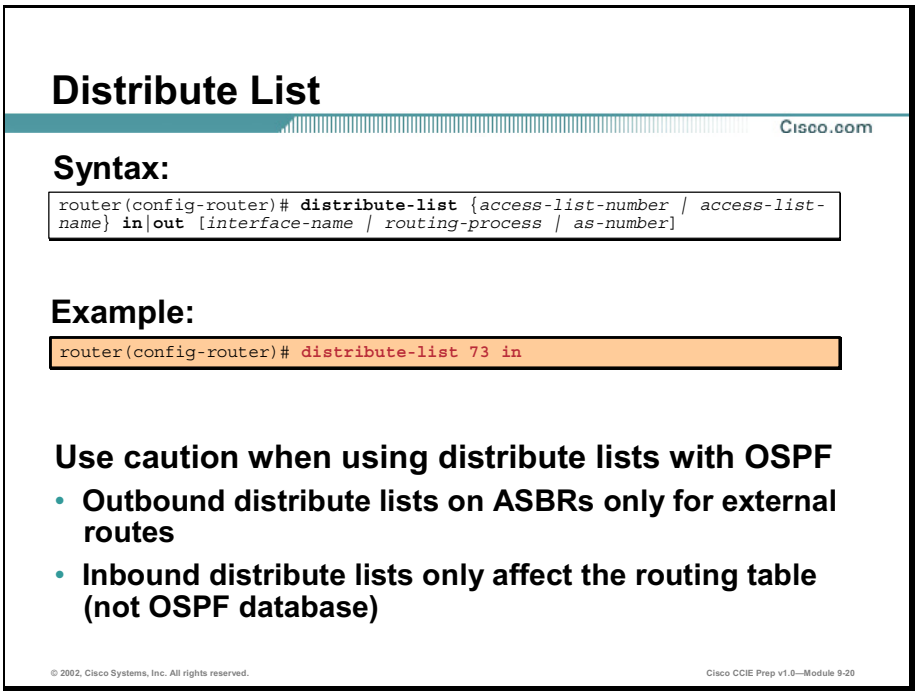

You might want to avoid processing certain routes listed in incoming updates. To suppress routes in incoming updates, use the following command in router configuration mode:

```
router(config-router)# distribute-list {access-list-number | access-list-name}
in|out [interface-name |routing-process | as-number]
```

Filtering sources of routing information prioritizes routing information from different sources, because some pieces of routing information may be more accurate than others. An *administrative distance* is a rating of the trustworthiness of a routing information source, such as an individual router or a group of routers. In a large network, some routing protocols and some routers can be more reliable than others as sources of routing information. Also, when multiple routing processes are running in the same router for IP, it is possible for the same route to be advertised by more than one routing process. By specifying administrative distance values, you enable the router to intelligently discriminate between sources of routing information. The router will always pick the route whose routing protocol has the lowest administrative distance.

With OSPF, routes cannot be filtered from entering the OSPF database. The **distribute-list in** command only filters routes from entering the routing table, but it does not prevent link-state packets from being propagated.

With OSPF the command **distribute-list out** works only on the routes being redistributed by the Autonomous System Boundary Routers (ASBRs) into OSPF. It can be applied to external type 2 and external type 1 routes, but not to intra-area and inter-area routes.

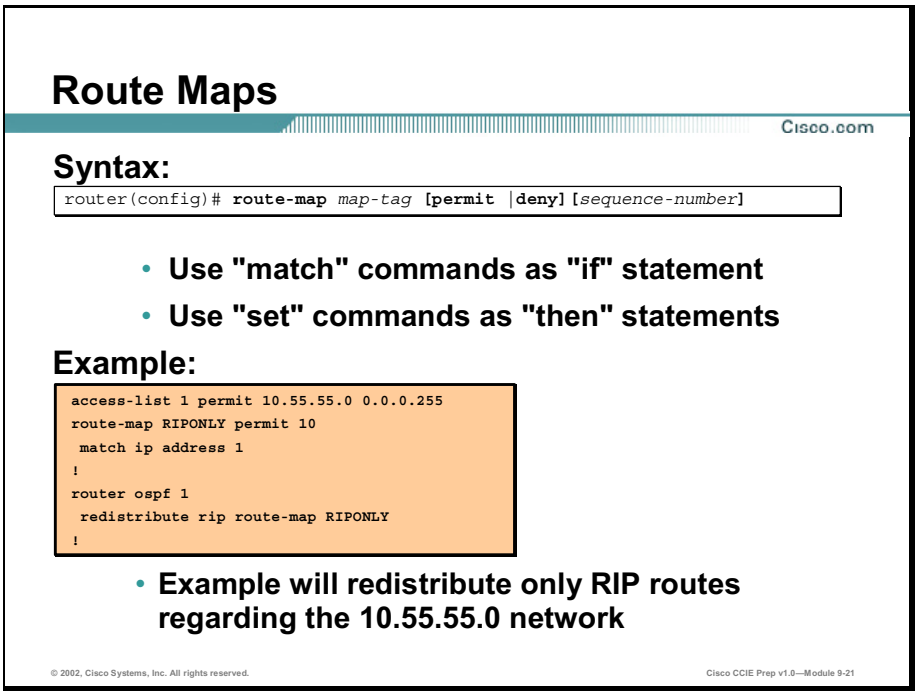

Route maps can be used to control routes when redistributing from one routing protocol to another. Route Maps can also be used for many functions, such as policy routing.

To define the conditions for redistributing routes from one routing protocol into another, use the **route-map** global configuration command and the **match** and **set** route-map configuration commands. Route Maps function much like “if” “then” statements in programming languages. To delete an entry, use the **no** form of this command. Route Maps can be used for many other purposes as well. For instance BGP or OSPF can use a route-map on the default-information originate command, to match a route in the route table, and based on this match, the routing protocol will either forward or stop the default route from being propagated.

- **route-map** *map-tag* [**permit** | **deny**] [*sequence-number*]
  - match length *min max*

    and/or
  - match ip address {*access-list-number* | *name*} [*...access-list-number* | *name*]
  - set ip precedence [*number* | *name*]
  - set ip next-hop *ip-address* [*... ip-address*]
  - set interface *interface-type interface-number* [*... type number*]

- set ip default next-hop *ip-address* [... *ip-address*]

- set default interface *interface-type*
  interface-number [... *type* ...*number*]

Use the match and set route-map configuration commands, to define the conditions for redistributing routes from one routing protocol into another. Each route-map command has a list of match and set commands associated with it. The match commands specify the *match criteria*—the conditions under which redistribution is allowed for the current route-map command. The set commands specify the *set actions*—the particular redistribution actions to perform if the criteria enforced by the match commands are met.

The match commands can be given in any order, and all match commands must "pass" to cause the route to be redistributed. Any route that does not match at least one match will be ignored; that is, the route will not be advertised.

The following example redistributes RIP routes that match the access list that is referenced from the route-map. Only the network 10.55.55.0 will be allowed into OSPF, regardless of how many networks are inside of the RIP process.

Example

```
access-list 1 permit 10.55.55.0 0.0.0.255
route-map RIPONLY permit 10
 match ip address 1
!
router ospf 1
redistribute rip route-map RIPONLY
```


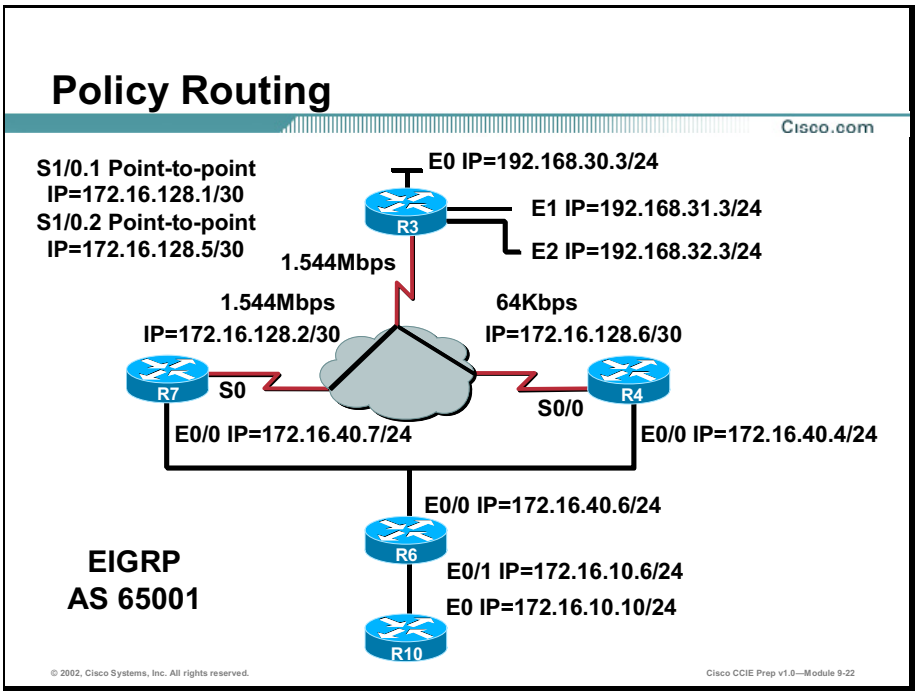

Policy Routing

Policy Routing provides the following benefits:

- **Source-Based Transit Provider Selection:** Policy-based routing to route traffic originating from different sources through different Internet connections across the policy routers.
- **Quality of Service (QoS):** QoS can be deployed by differentiating traffic by setting the precedence or Type of Service (ToS) values in the IP packet headers at the edge of the network.
- **Cost Savings:** Cost savings can be achieved by distributing interactive and batch traffic among low-bandwidth, low-cost permanent paths and high-bandwidth, high-cost paths.
- **Load Sharing:** Implement policies to distribute traffic among multiple paths based on the traffic characteristics.

Policy-based routing is applied to incoming packets. All packets received on an interface with policy-based routing enabled are considered for policy-based routing. The router passes the packets through route maps and makes decisions based on the criteria defined in the route maps. Packets are then forwarded/routed to the appropriate next hop based on the route-map.

Note	Policy routing is specified on the interface that receives the packets, not on the interface from which the packets are sent.

Policy Routing is configured with the interface command:

```
router(config-if)# ip policy route-map route-map_name
```

To use policy routing on packets sourced from the router use the following global configuration command:

```
router(config)# ip local policy route-map route-map_name
```

Example

In the example, R10 has two paths to the networks 192.168.31.0/24 and 192.168.30.0/24. EIGRP is the routing protocol and will traffic share any traffic coming from R6 destined towards 192.168.31.0/24 and 192.168.30.0/24 between R7 and R4. In this example we will define a policy route on R6, E0/1, that states any IP traffic for 192.168.31.0/24 and 192.168.30.0/24 will always go through R4.

Here is an example of a traceroute illustrating the use of 172.16.40.7 for the second hop.

```
r10# trace 192.168.31.3

Type escape sequence to abort.
Tracing the route to 192.168.31.3

  1 172.16.10.6 4 msec 0 msec 4 msec
  2 172.16.40.7 4 msec 0 msec 4 msec
  3 172.16.128.1 16 msec 4 msec *
r10#
```

This example lists the configuration of R6.

```
hostname r6
!!
interface Ethernet0/0
 ip address 172.16.40.6 255.255.255.0
!
interface Ethernet0/1
 ip address 172.16.10.6 255.255.255.0
 ip policy route-map take_r4

!
access-list 101 permit ip any 192.168.30.0 0.0.1.255
route-map take_r4 permit 10
 match ip address 101
```

```
 set ip next-hop 172.16.40.4
```

The following example shows a trace on R10 after the policy route was configured.

```
r10# trace 192.168.31.3

Type escape sequence to abort.
Tracing the route to 192.168.31.3

  1 172.16.10.6 4 msec 0 msec 4 msec
  2 172.16.40.4 4 msec 4 msec 4 msec
  3 172.16.128.5 20 msec 12 msec *
r10#
```

Summary

This section summarizes the key points discussed in this lesson.

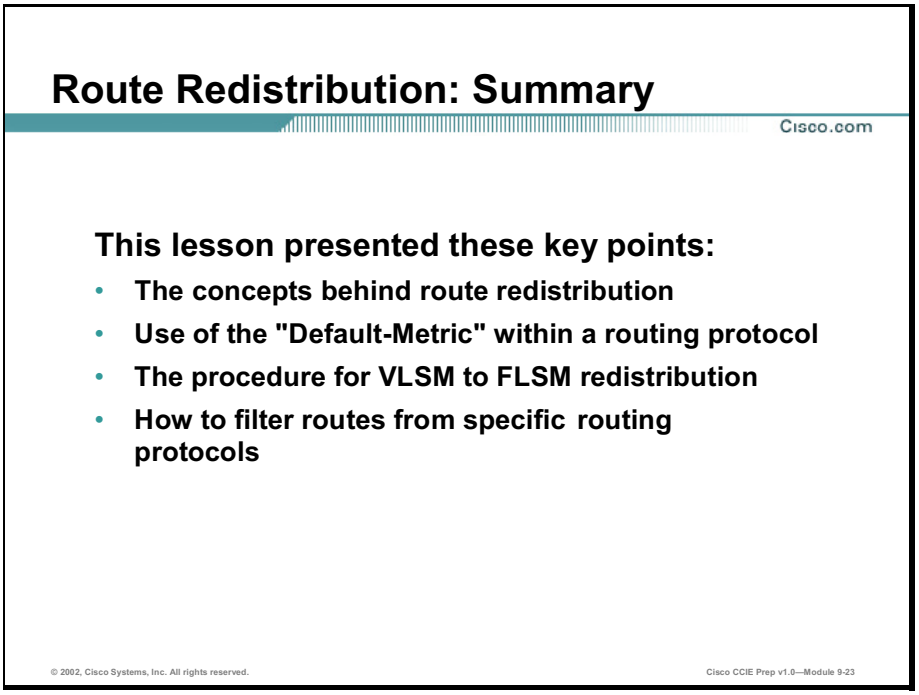

Next Steps

After completing this lesson, go to:

- Authentication

References

For additional information, refer to these resources:

- http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fipr_c/ipcprt2/1cfindep.htm

Lesson Assessment (Quiz)

Q1) Using EIGRP, you notice that your subnets do not show up across the entire network. What can you do to correct this?

A) Manually redistribute from EIGRP into OSPF, modify the summary address, then redistribute back into EIGRP

B) Use the *subnets* option for redistribution

C) Use the *no auto-summarize* option

D) This situation cannot be corrected with today's technology

Q2) What are the safe techniques for redistribution of routes, without creating a routing loop?

A) Avoid mutual redistribution

B) Use route maps to only allow specific routes in the redistribution

C) Designate OSPF over ISDN as demand circuits

D) Use snapshot routing

Q3) On an ASBR you use the **area range** command, but the redistributed RIP routes are not being summarized into OSPF. What would cause this?

A) The **area range** command only works on classful boundaries

B) The *subnets* option should be removed within the redistribution statement

C) The **area range** command only summarized OSPF routes, no redistributed routes

D) OSPF can support VLSM, but routes redistributed from RIP must all use the same mask forever

Q4) How can you redistribute a 28-bit OSPF route into a 26-bit RIPv1 domain?

A) Create static routes with a 26-bit mask that correlate to the OSPF routes, and redistribute those into RIP

B) Allow OSPF to summarize the 28-bit mask networks into a 26-bit mask using the **area range** command

C) Redistribute the OSPF routes into EIGRP, and allow EIGRP to summarize the routes to a 26-bit route on an interface-by-interface basis

D) Use the **redistribute** command, with the *subnets* option

Authentication

Overview

When routers exchange updates via a routing protocol, there is no authentication by default. In many networks, you want to add the authentication so that a rogue router is not advertising its routes within the routing domain. In this lesson you will learn how, when configured, neighbor authentication occurs whenever routing updates are exchanged between neighbor routers.

Importance

Authentication ensures that a router receives reliable routing information from a trusted source.

Objectives

Upon completing this lesson, you will be able to configure authentication on the following protocols:

- RIPv2
- EIGRP
- OSPF
- IS-IS
- BGP

Learner Skills and Knowledge

To fully benefit from this lesson, you must have these prerequisite skills and knowledge:

- Passed the Cisco Certified Internetworking Expert (CCIE) written qualification exam
- Completed the Cisco course Building Scalable Cisco Internetworks (BSCI) or knowledge of Link State and Distance-Vector protocols, their operations, and how to configure them

Outline

This lesson includes these sections:

- Overview
- Authentication Concepts
- OSPF Authentication
- RIPv2 Authentication
- IS-IS Authentication
- EIGRP Authentication
- BGP Authentication
- Summary
- Lesson Assessment (Quiz)

Authentication Concepts

The authentication ensures that a router receives reliable routing information from a trusted source. This section examines authentication concepts.

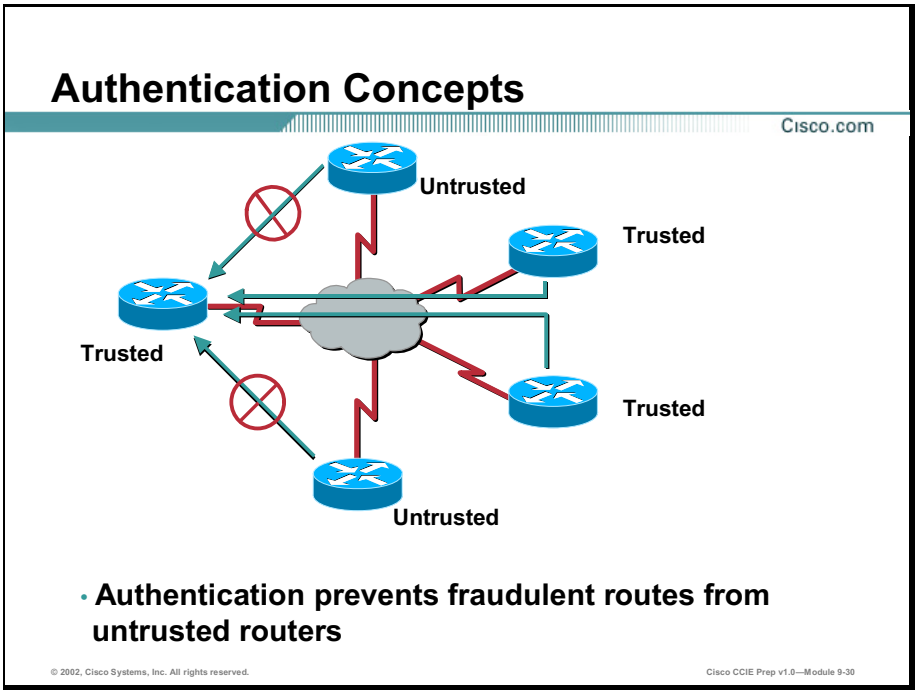

Without neighbor authentication, unauthorized or deliberately, malicious routing updates could compromise the security of your network traffic. A security compromise could occur if an unfriendly party diverts or analyzes your network traffic. For example, an unauthorized router could send a fictitious routing update to convince your router to send traffic to an incorrect destination. This diverted traffic could be analyzed to learn confidential information of your organization, or merely used to disrupt your organization's ability to effectively communicate using the network. Neighbor authentication prevents any such fraudulent route updates from being received by your router.

Neighbor authentication can be configured for the following routing protocols:

- Border Gateway Protocol (BGP)
- Director Response Protocol (DRP) Server Agent
- Intermediate System-to-Intermediate System (IS-IS)
- IP Enhanced Interior Gateway Routing Protocol (EIGRP)
- Open Shortest Path First (OSPF)
- Routing Information Protocol (RIP) version 2

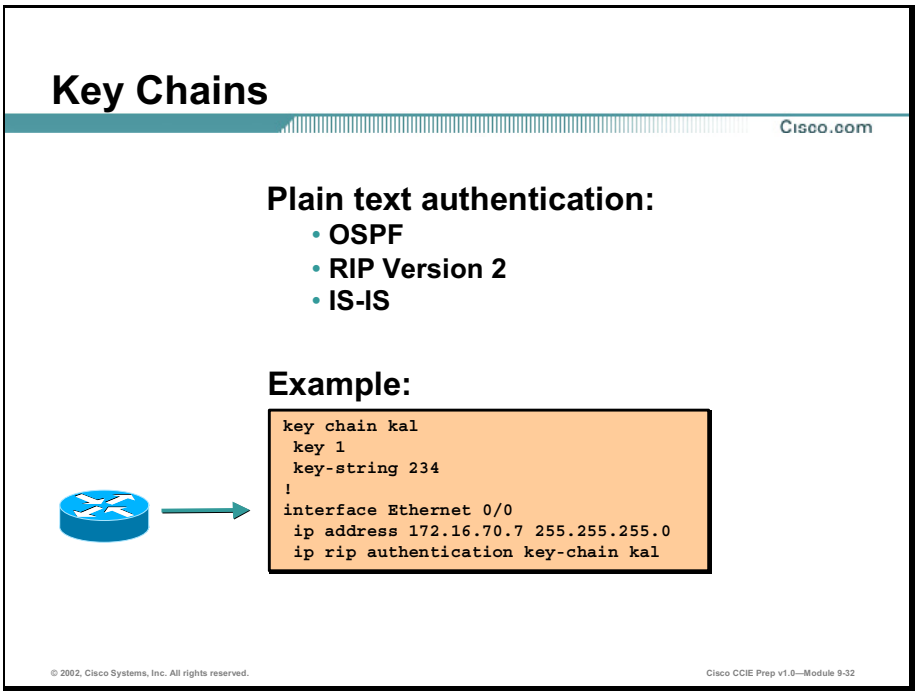

You can configure key chains for these IP routing protocols:

- Routing Information Protocol v2 (RIPv2)
- IP Enhanced Interior Gateway Routing Protocol (IGRP) (supports Message Digest Version 5 (MD5) authentication only)
- DRP Server Agent

These routing protocols offer the additional function of managing keys by using key chains. When you configure a key chain, you specify a series of keys with lifetimes, and the Cisco Internetwork Operating System (IOS) software rotates through each of these keys. This decreases the likelihood that keys will be compromised.

Each key definition within the key chain must specify a time interval for which that key will be activated (its "lifetime"). Then, during a given key's lifetime, routing update packets are sent with this activated key. Keys cannot be used during time periods for which they are not activated. Therefore, it is recommended that for a given key chain, key activation times overlap to avoid any period of time for which no key is activated. If a time period occurs during which no key is activated, neighbor authentication cannot occur, and therefore routing updates will fail.

Multiple key chains can be specified.

Note	The router needs to know the time to be able to rotate through keys in synchronization with the other participating routers, so that all routers are using the same key at the same moment.

Key Chain Example

```
key chain kal
 key 1
 key-string 234
!
interface Serial2
ip address 141.108.0.10 255.255.255.252
 ip rip authentication key-chain kal
```

OSPF Authentication

This section covers how to configure plain text as well as MD5 authentication on a Cisco router running Open Shortest Path First (OSPF).

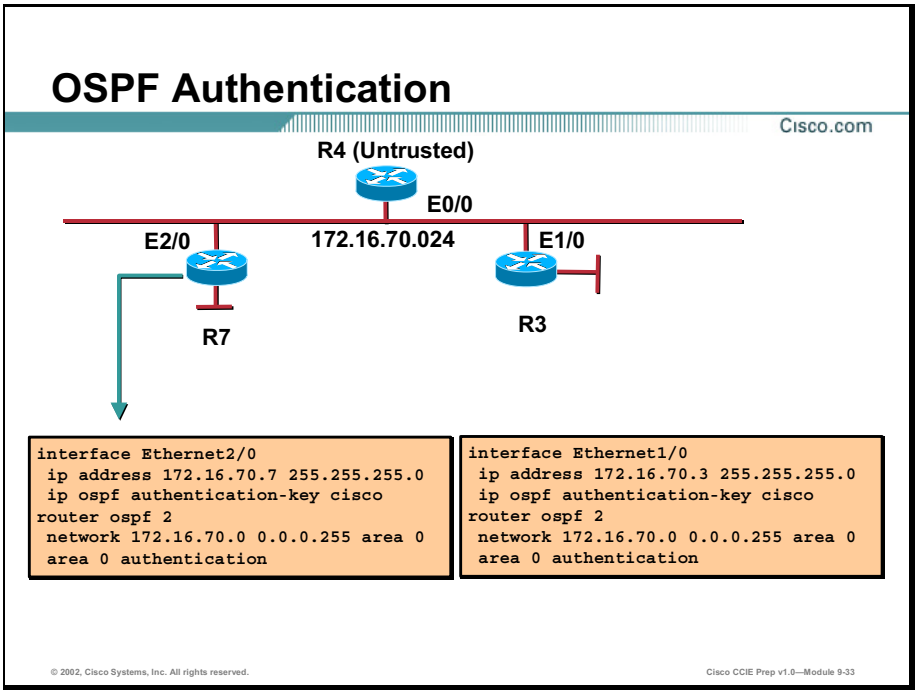

Assign a password to be used by neighboring OSPF routers on a network segment that is using the OSPF simple password authentication.

ip ospf authentication-key *key*

OSPF Plain Text

```
interface Ethernet2/0
 ip address 8.0.0.1 255.0.0.0
 ip ospf authentication-key cisco
router ospf 2
 network 8.0.0.0 0.255.255.255 area 0
 area 0 authentication
```

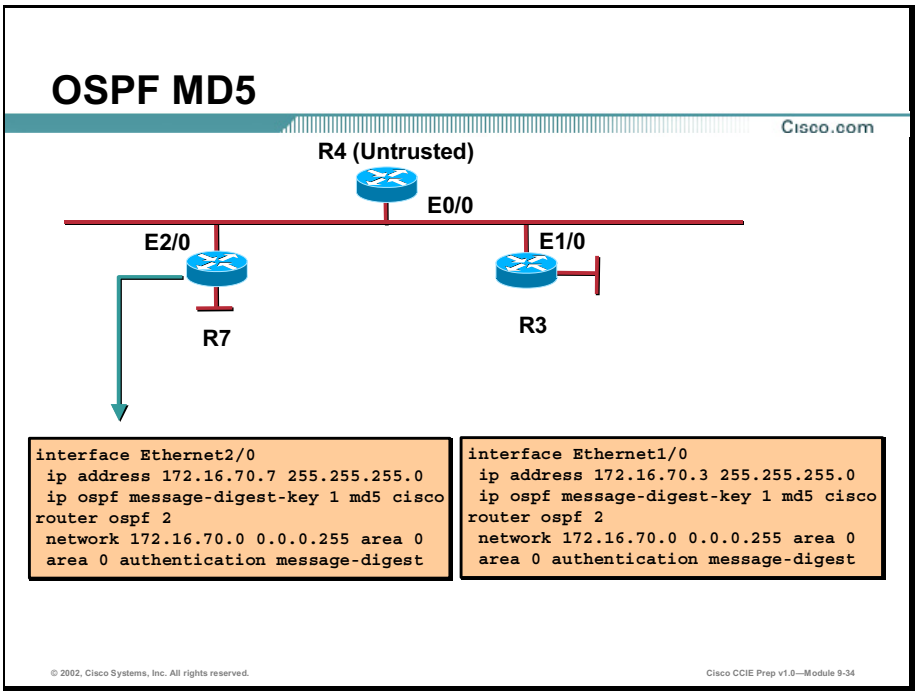

Enable OSPF MD5 authentication. The values for *keyid* and *key* must match values specified for other neighbors on a network segment.

ip ospf message-digest-key *keyid* **md5** *key*

OSPF MD5

```
interface Ethernet2/0
 ip address 8.0.0.1 255.0.0.0
 ip ospf message-digest-key 1 md5 cisco

router ospf 2
 network 8.0.0.0 0.255.255.255 area 0
 area 0 authentication message-digest
```

To verify authentication, you can use **debug ip ospf adj**. If you clear all routes, and get your OSPF routes back, that would also allow you to verify that it is working.

Note	Authenticating Area 0 and Virtual Links: If a virtual link is configured, and authentication is enabled for Area 0, both ends, or the routers at each end of the virtual link must also be configured for authentication. Remember that a virtual link is an extension of area 0. Anything that is configured in area 0 must also be configured on the router at the other end of the virtual link.

RIPv2 Authentication

This section explains how to configure authentication on a Cisco router running RIPv2.

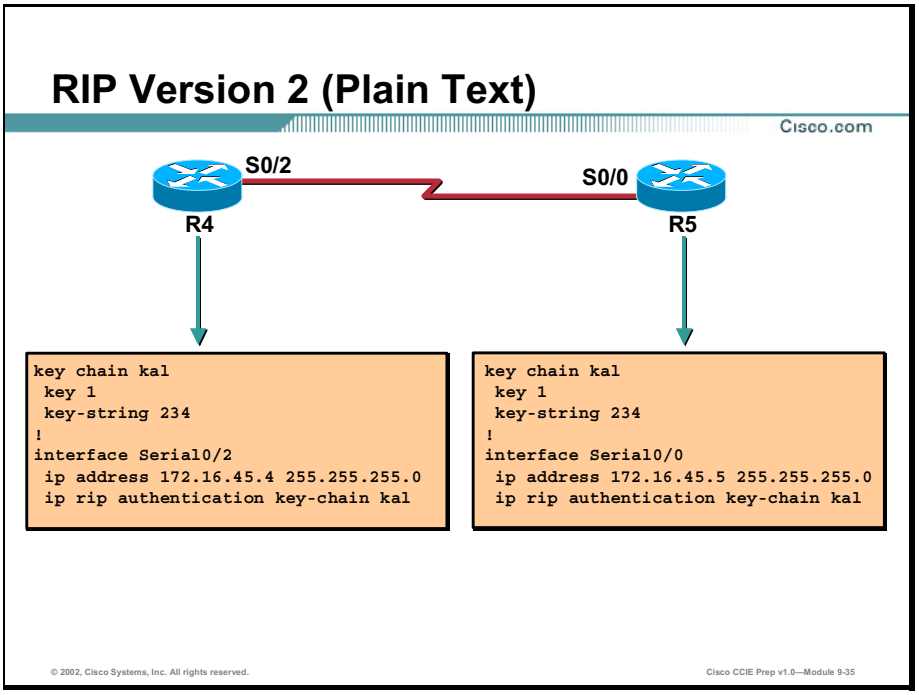

Routing Information Protocol (RIP) only supports authentication when RIPv2 is running. To enable authentication, you need to define a key chain and then apply that key chain to an interface.

RIPv2, Plain Text

```
Hostname Router1
!
key chain kal
  key 1
  key-string 234
!
interface Serial0/2
  ip address 172.16.45.4 255.255.255.0
  ip rip authentication key-chain kal
```

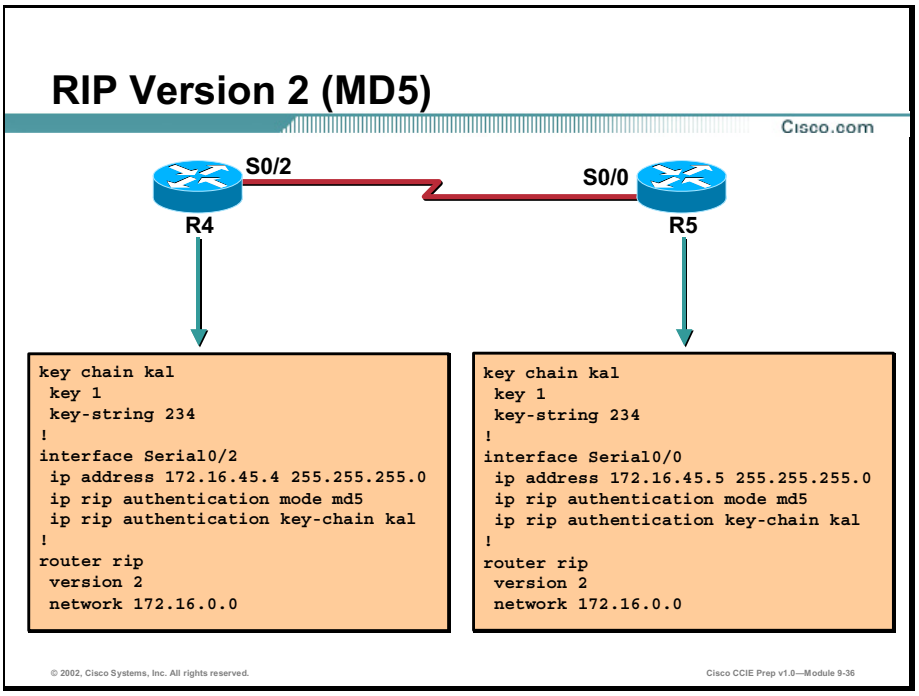

The following is an example of RIPv2 authentication using MD5.

RIPv2, MD5

```
key chain kal
 key 1
 key-string 234
!
interface Serial0/2
 ip address 172.16.45.4 255.255.255.0
 ip rip authentication mode md5
 ip rip authentication key-chain kal
!
router rip
 version 2
 network 172.16.0.0
```

IS-IS Authentication

This section covers how to configure authentication on a Cisco router running Intermediate System to Intermediate System (IS-IS).

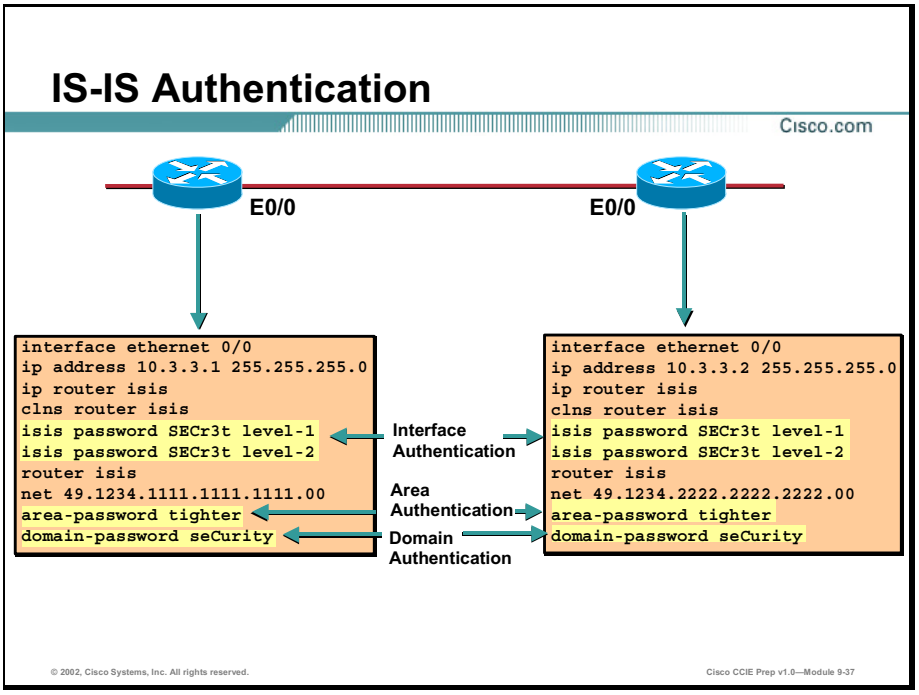

This is an example of authentication between IS-IS neighbors. Notice that there is a level-1 and a level-2 password defined in the example. There is also an area and a domain password defined.

IS-IS Authentication Plain Text

```
Interface authentication
interface ethernet 0/0
 ip address 10.3.3.1 255.255.255.0
 ip router isis
 clns router isis
 isis password SECr3t level-1
 isis password SECr3t level-2

Area authentication

router isis
 net 49.1234.1111.1111.1111.00
 area-password tighter
Domain Authentication
```

```
router isis
 net 49.1234.1111.1111.1111.00
 domain-password seCurity
```

EIGRP Authentication

This section describes how to configure authentication on a Cisco router running Enhanced Interior Gateway Routing Protocol (EIGRP).

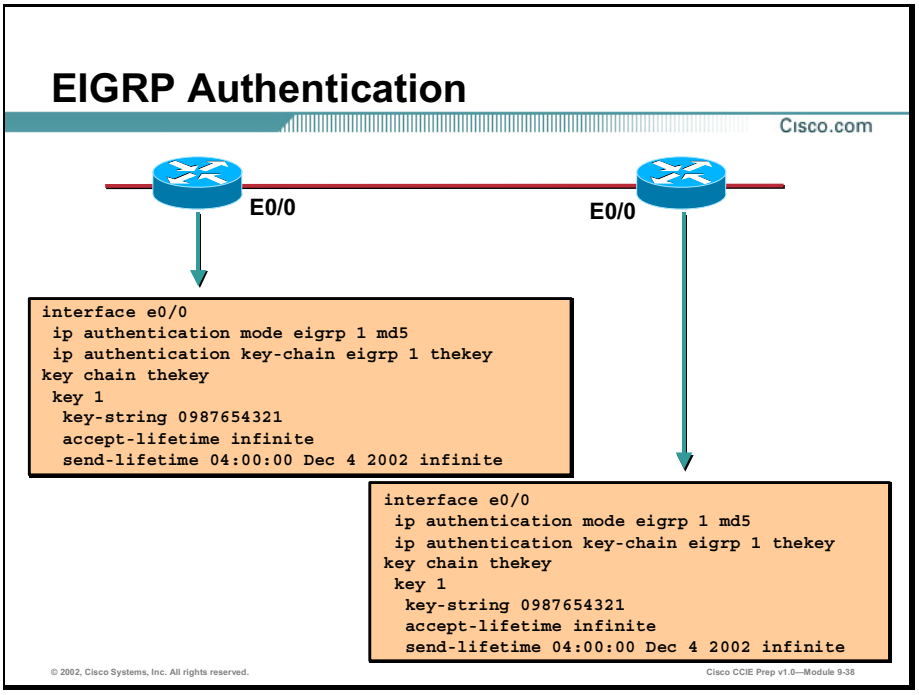

This is an example of authentication between EIGRP neighbors. EIGRP authentication is very similar to RIPv2 authentication in that you define a key chain and apply that key chain to an interface.

EIGRP Authentication

```
Interface e0/0
   ip authentication mode eigrp 1 md5
   ip authentication key-chain eigrp 1 thekey
key chain thekey
  key 1
    key-string 0987654321
    accept-lifetime infinite
    send-lifetime 04:00:00 Dec 4 2002 infinite
```

BGP Authentication

This section describes how to configure authentication on a Cisco router running Border Gateway Protocol (BGP).

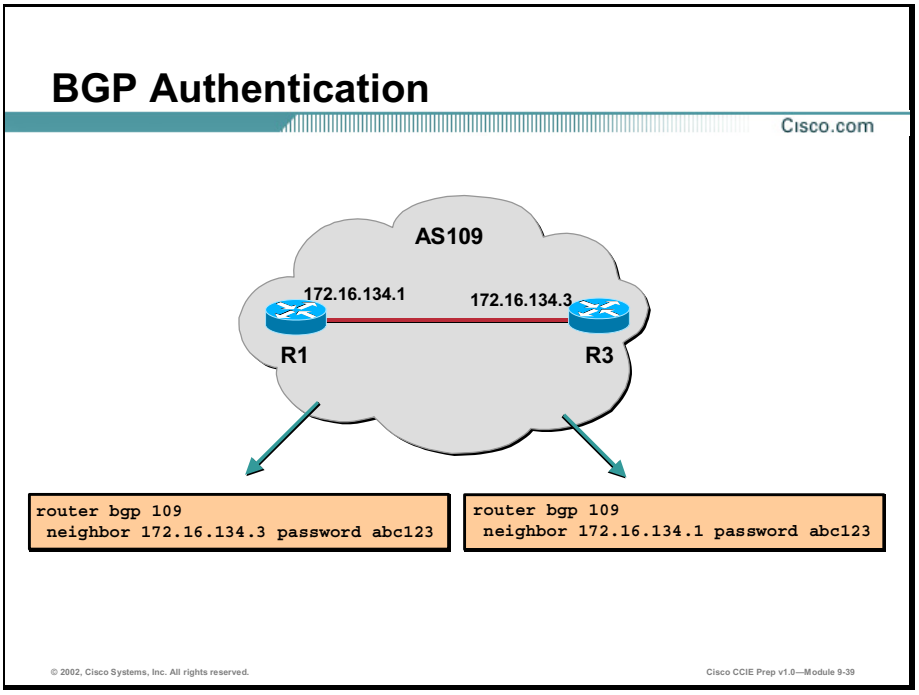

The following example specifies that the router and its BGP peer at 172.16.134.1 invoke MD5 authentication on the Transfer Control Protocol (TCP) connection that is between them.

TCP MD5 Authentication for BGP Example

```
router bgp 109
 neighbor 172.16.134.3 password abc123
```

Summary

This section summarizes the key points discussed in this lesson.

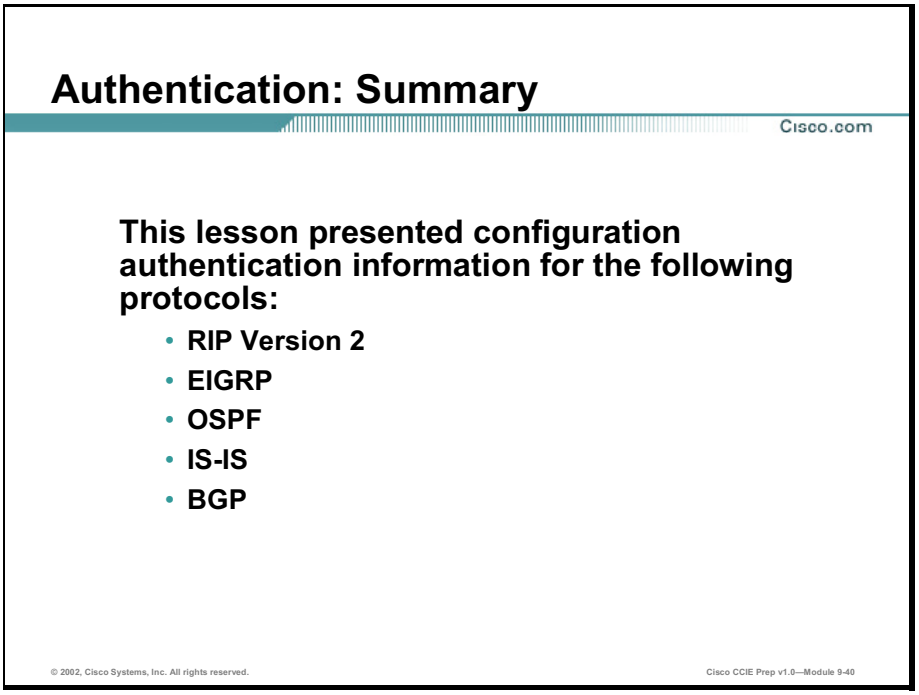

Next Steps

After completing this lesson, go to:

- Desktop Protocols

References

For additional information, refer to these resources:

- http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fipr_c/ipcprt2/1cfindep.htm

Lesson Assessment (Quiz)

Q1) On the network, some of the routers receive RIP routes and others do not. What could cause this?

A) A router may be directly connected to all networks

B) Distribute lists may be applied

C) The version of RIP may not be matched either globally or on an interface-by-interface basis

D) Authentication may be set incorrectly on some of the routers

E) The passive interface option may be prohibiting some of the routers from receiving updates

9

Bridging and DLSW+

Overview

In addition to Internet Protocol (IP), the Cisco Certified Internetworking Expert (CCIE) Lab candidate should also be proficient with various bridging techniques and the transportation of non-routable protocols. This module will examine bridging techniques such as Transparent Bridging and Integrated Routing and Bridging (IRB). This module will also address the configuration and troubleshooting of Data Link Switching Plus (DLSw+).

Upon completing this module, you will be able to:

- Describe the need for bridging, and configure Cisco's supported bridging technologies
- Configure DLSw+ features
- Isolate and resolve DLSw+ problems on the network

Outline

The module contains these lessons:

- Bridging Concepts
- Data Link Switching Plus Concepts
- Troubleshooting DLSw+

Bridging Concepts

Overview

This lesson will cover commonly used bridging technologies used in Ethernet environments.

Importance

A CCIE Lab candidate needs to know how to configure Cisco routers to perform Transparent Bridging and Integrated Routing and Bridging (IRB).

Objectives

Upon completing this lesson, you will be able to:

- Describe the operation of Transparent Bridging
- Explain the purpose of Integrated Routing and Bridging

Learner Skills and Knowledge

To fully benefit from this lesson, you must have these prerequisite skills and knowledge:

- Passed the CCIE written qualification exam
- An understanding of the fundamental concepts of the Open Systems Interconnection (OSI) model, including Layer 2 and Layer 3 functions

Outline

This lesson includes these sections:

- Overview
- Bridging Overview
- Transparent Bridging
- Integrated Routing and Bridging
- Summary
- Lesson Assessment (Quiz)

Bridging Overview

This section provides an overview of the function of bridges and routers in an internetwork.

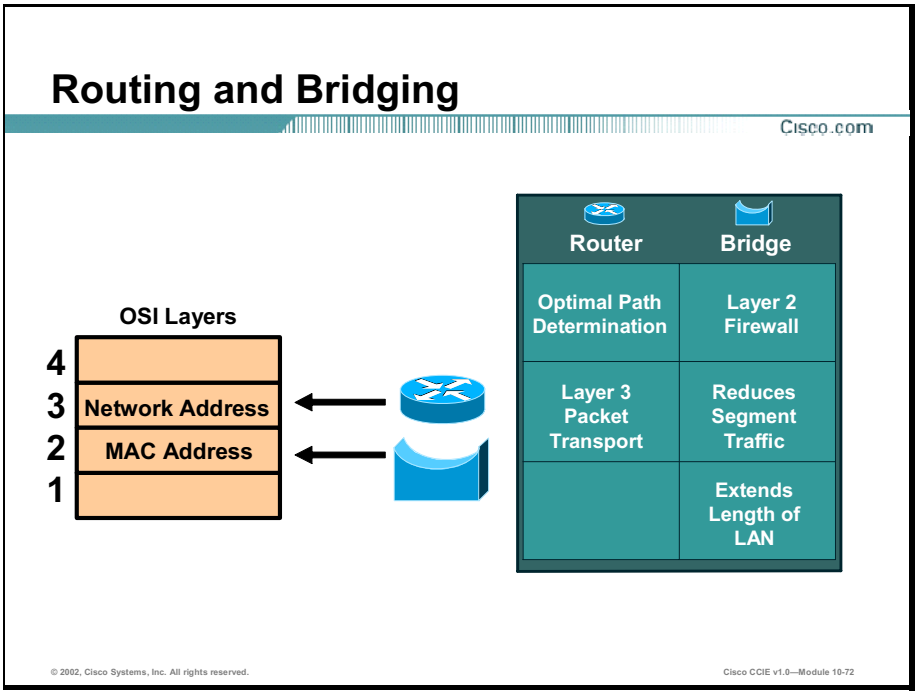

Bridges are internetworking devices designed to interconnect Local Area Networks (LANs) to form the appearance of a single larger data link. Bridges analyze incoming frames, make forwarding decisions based on information in the frames, and forward the frames to their destination. A bridge isolates traffic, acts as a Layer 2 firewall, reduces traffic on any given segment, and extends the effective length of the LAN.

The disadvantages of bridges are unthrottled broadcasts, session timeouts, and hop-count limitations.

Routers move information across an internetwork from source to destination. Routing involves two activities: determining the optimal path, and transporting the information (packets) to their destination. Routers can select an optimal path; bridges do not select optimal paths. Bridging works at the data link layer, whereas routing works at the network layer.

All devices on bridged segments are part of the same network address space, whereas routing allows networks to be divided into smaller units.

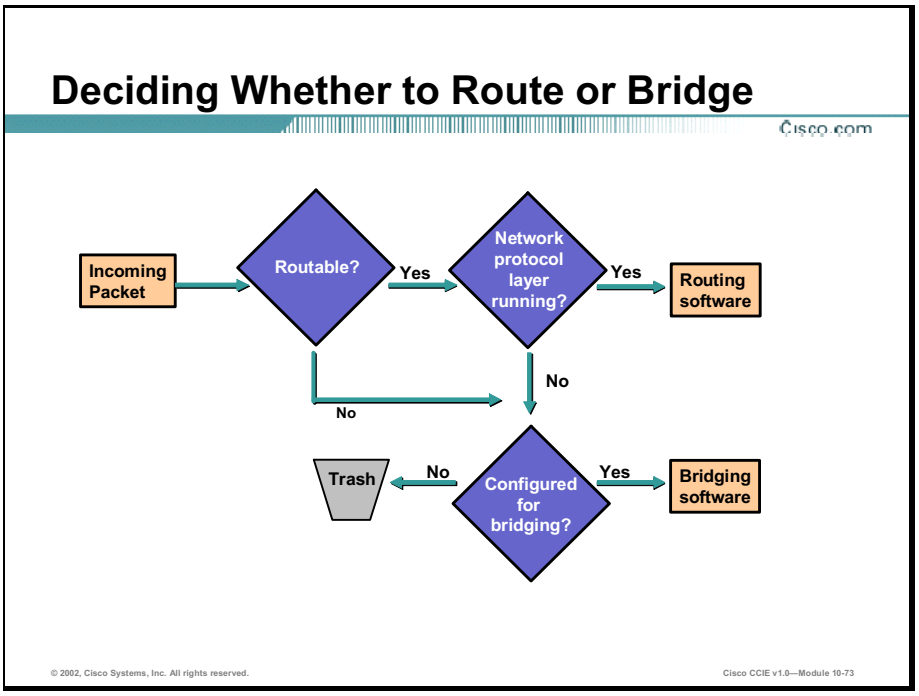

Packets arrive at the router. If the protocol has a network layer, or if the network layer protocol is operating in the router, then the packet is passed to the routing software.

If the protocol has no network layer or the network layer protocol is not configured in the router, then the packet is passed to the bridging software if the router is configured for bridging.

With integrated routing and bridging, certain routable protocols can be bridged across bridged groups and the same protocol can be routed through a routed interface. You will learn about integrated routing and bridging in a later session.

Note	If the router does not have routing enabled for a particular layer three protocol, such as Internetwork Packet Exchange (IPX), or AppleTalk, the router will bridge those frames when it receives them.

Transparent Bridging

This section covers the tasks necessary to configure transparent bridging.

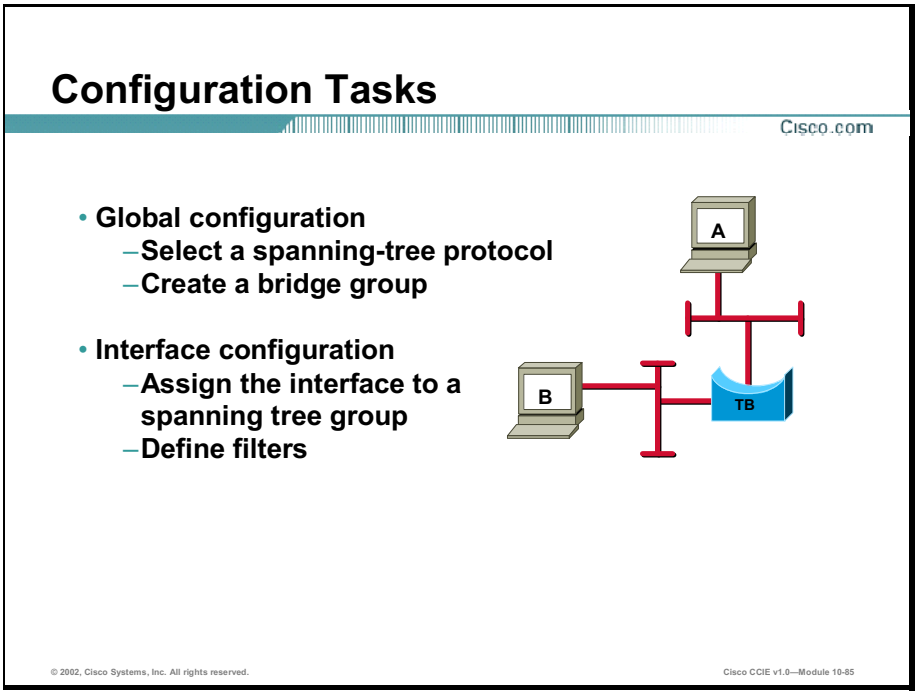

This slide provides an overview of the configuration tasks required to implement transparent bridging and the spanning-tree algorithm. Each of these commands is presented in detail in the next two slides.

The only required configuration tasks are:

- Selecting a Spanning-Tree Protocol
- Assigning the interface to a bridge group

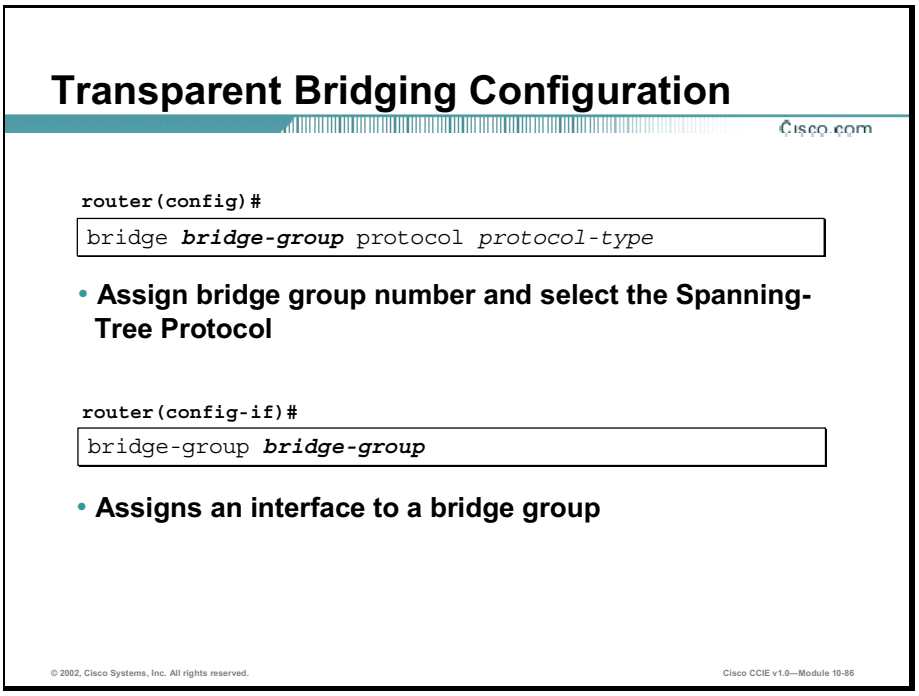

The bridge-group number is a number ranging from 1 - 255, protocol-type is either Digital Equipment Corporation (DEC) or Institute of Electrical and Electronics Engineers (IEEE). An interface is then associated with a bridge group with the command **bridge-group** ***bridge-group***.

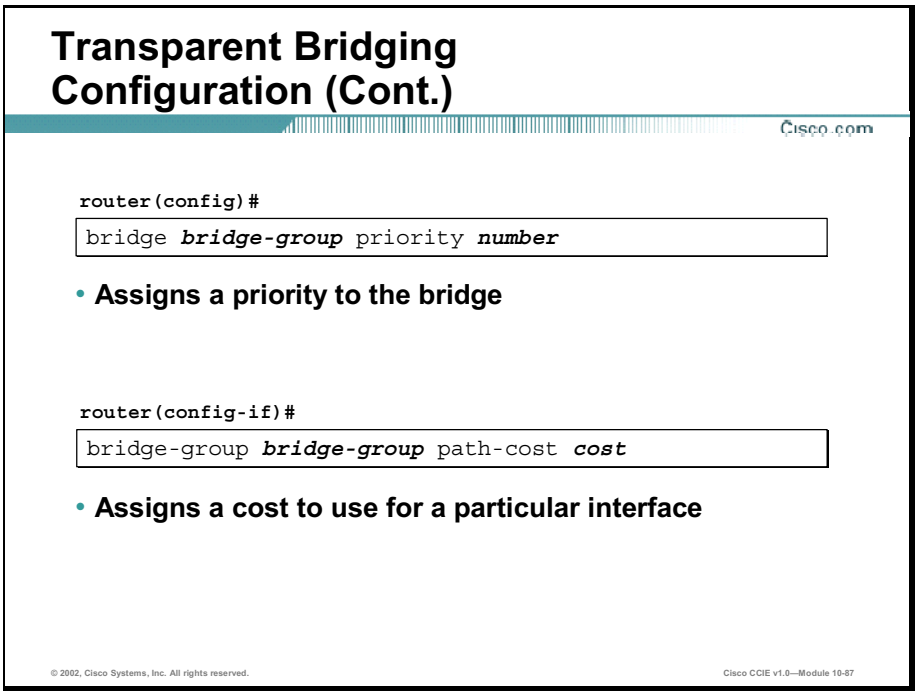

The priority number can range from 0 to 255 (DEC), or 0 to 64000 (IEEE).

- The default priority for IEEE is 32768.
- The default priority for DEC is 128.

Path cost can range from 1 to 65535, with higher values indicating higher costs. This range applies regardless of whether the IEEE or DEC Spanning Tree Protocol has been specified.

By convention, the default path cost is 10000/data rate of the attached LAN (IEEE), or 100000/data rate of the attached LAN (DEC), in megabits per second.

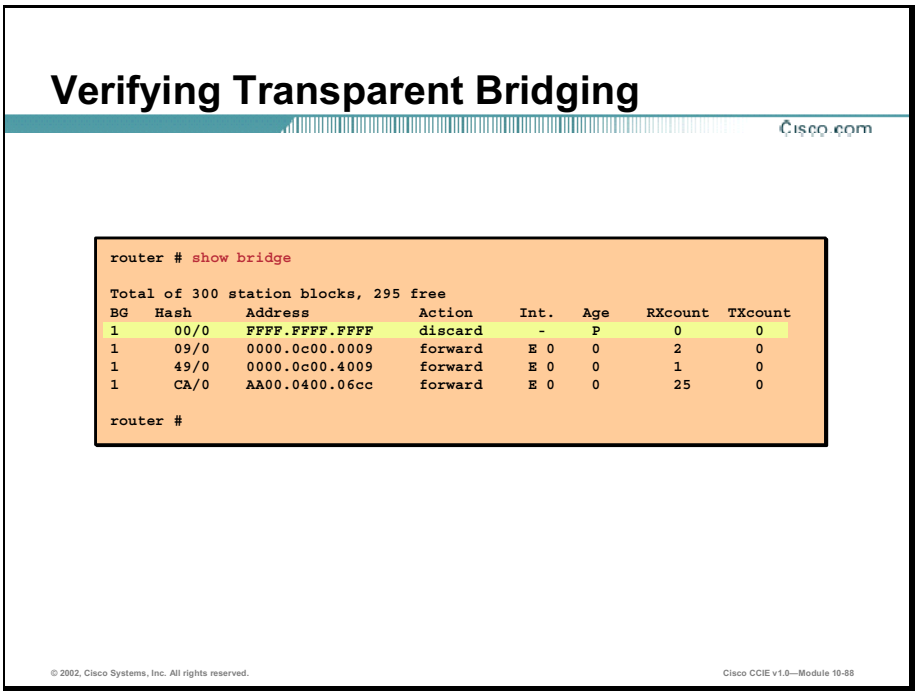

This list describes the output from the **show bridge** command shown:

- **Total of 300 station blocks:** Total number of forwarding database elements in the system. The memory to hold bridge entries is allocated in blocks of memory sufficient to hold 300 individual entries. When the number of free entries falls below 25, another block of memory sufficient to hold another 300 entries is allocated. Therefore, the size of the bridge-forwarding database is limited to the amount of free memory in the router.
- **295 free:** Number in the free list of forwarding database elements in the system. The total number of forwarding elements is expanded dynamically, as needed.
- **BG:** Bridging group to which the address belongs.
- **Hash:** Hash key/relative position in the keyed list.
- **Address:** Canonical (Ethernet ordered) Media Access Control (MAC) address.
- **Action:** Action to be taken when that address is looked up; choices are to discard or forward the datagram.
- **Interface:** Interface, if any, on which that address was seen.
- **Age:** Number of minutes since a frame was received from or sent to that address. The letter "P" indicates a permanent entry. The letter "S" indicates the system as recorded by the router. On the modular systems, this is typically the broadcast address and the router's own hardware address; on the IGS router, this field will also include certain multicast addresses.

- **RX count:** Number of frames received from that address.
- **TX count:** Number of frames forwarded to that address.

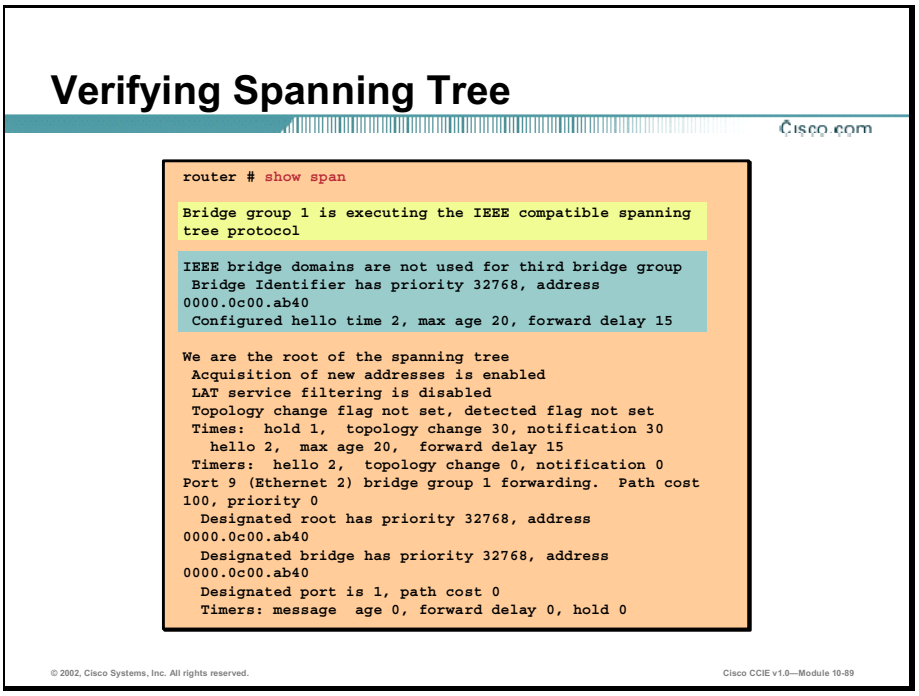

The output from the **show spanning-tree** command displays the type of Spanning Tree protocol running (IEEE compatible in this case), the bridge group identifier, and the bridge priority.

Notice this router (bridge) is the root bridge. This should be the only router with this indication.

Integrated Routing and Bridging

This section describes the need for and the configuration of Integrated Routing and Bridging (IRB).

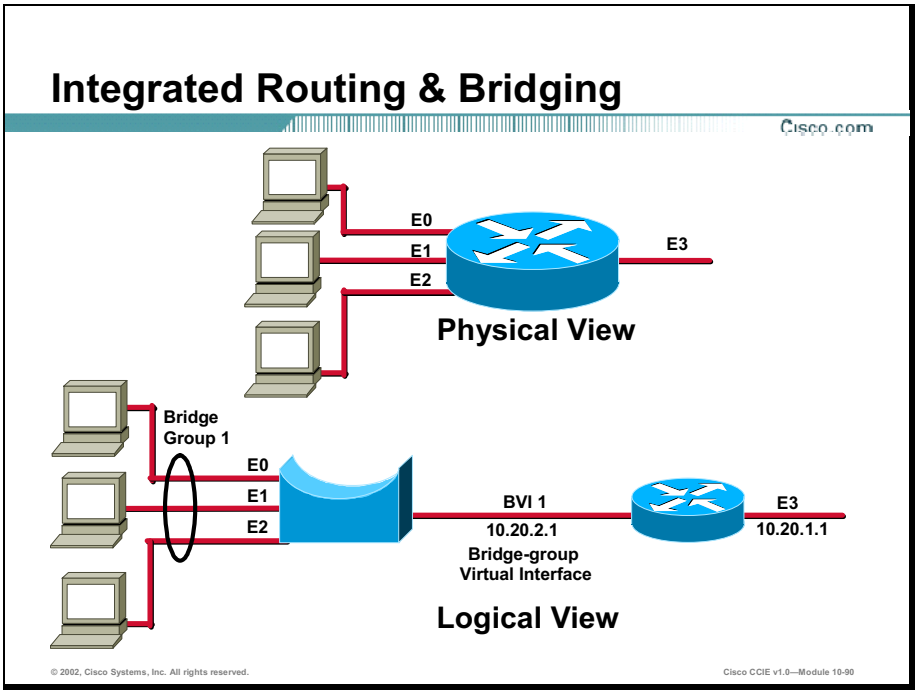

Integrated Routing and Bridging (IRB) can cause some interfaces to be bridging for a particular protocol (e.g., IP), while other interfaces can be routing for that protocol. This is accomplished via a Bridge-group Virtual Interface (BVI), which has both a protocol address and a bridge-group membership.

This example shows a router with four Ethernet interfaces. Three of the interfaces (E0 – E2) are bridged together in one bridge group and share a common data link. The three Ethernet interfaces are routed to other interfaces on the router. The configuration is shown in the following example.

Integrated Routing and Bridging Configuration

The three Ethernet interfaces (E0 – E2) are bound to the Bridge-group Virtual Interface (BVI) with an interface number corresponding to the respective bridge group.

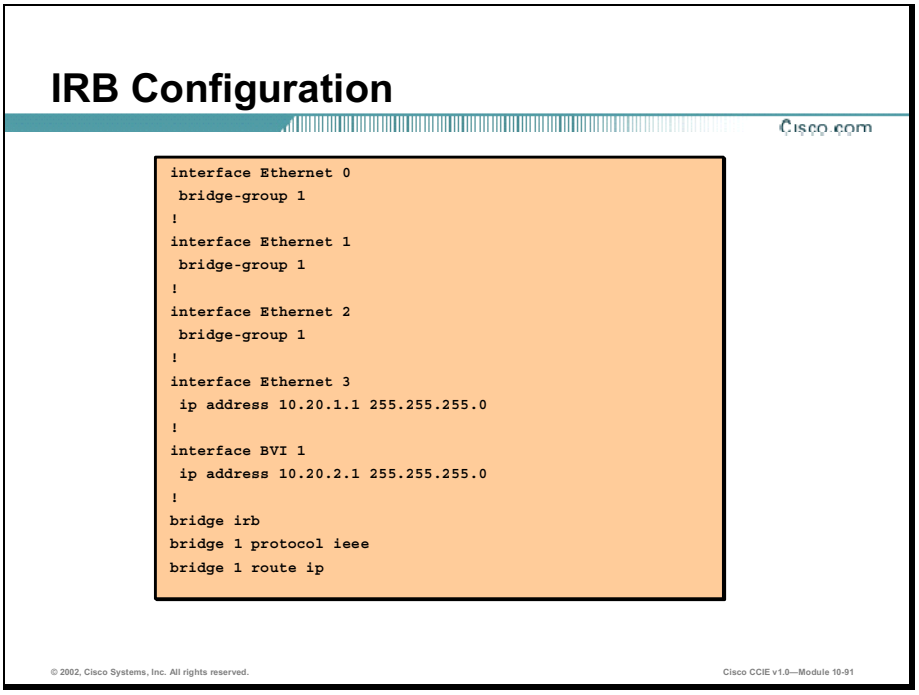

The **bridge irb** global configuration command enables Integrated Routing and Bridging. The **bridge 1 route ip** configuration command allows routing IP traffic from the bridge-group.

Summary

This section summarizes the key points discussed in this lesson.

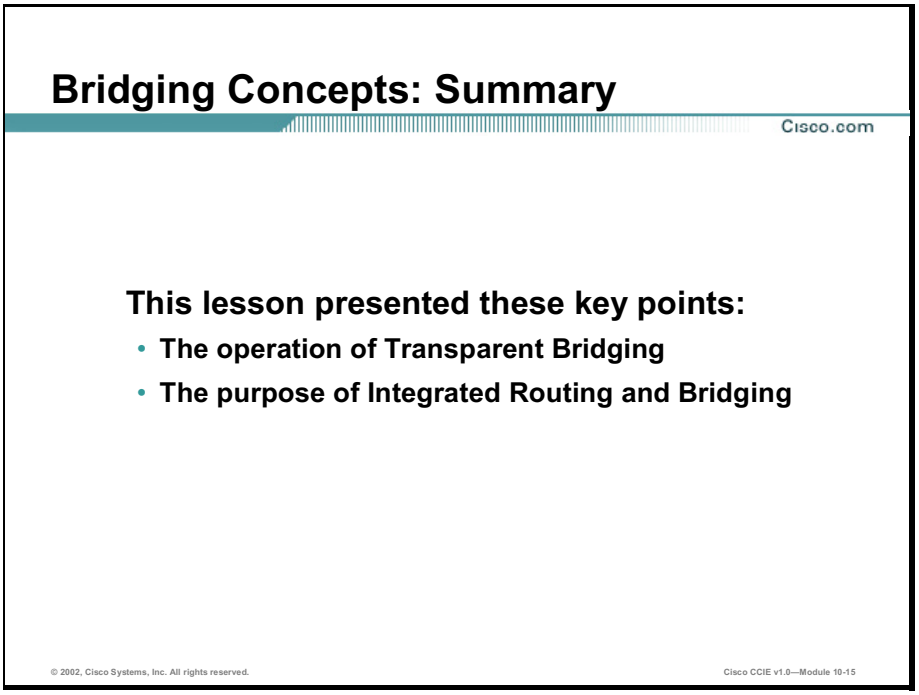

Next Steps

After completing this lesson, go to:

- Data Link Switching Plus Concepts

References

For additional information, refer to these resources:

- Transparent Bridging: http://www.cisco.com/warp/public/701/37.html
- Understanding and Configuring VLAN Routing and Bridging on a Router Using the IRB Feature: http://www.cisco.com/warp/public/473/741_10.html

Lesson Assessment (Quiz)

Q1) _____ are internetworking devices designed to interconnect LANs to form the appearance of a single larger data link.

A) Bridges

B) Firewalls

C) Routers

D) All of the above

Q2) Which of the following technologies allow a protocol to be bridged on one set of interfaces and routed on another set of interfaces?

A) SRB

B) SRT

C) IRB

D) SR/TLB

Q3) What is the default bridge priority when using the IEEE protocol?

A) 0

B) 32768

C) 128

D) 100

E) 16384

Data Link Switching Plus Concepts

Overview

Data Link Switching Plus (DLSw+) is one of the most efficient ways to transport non-routable data across the internetwork, such as time sensitive protocols like Systems Network Architecture (SNA). DLSw+ networks are very reliable and offer many options for scalability. This lesson will introduce the concepts surrounding DLSw+, including specific configuration examples.

Importance

DLSw+ is primarily used in the modern network to transport SNA. It has proven to be a very effective and robust protocol for handling time-sensitive data. DLSw+ is an integral component of the CCIE Lab and a significant percentage of lab points can be expected to come from DLSw+ configuration.

Objectives

Upon completing this lesson, you will be able to:

- Describe DLSw+ basic configuration
- Configure multiple encapsulation options
- Configure DLSw+ scalability features
- Configure advanced DLSw+ features

Learner Skills and Knowledge

To fully benefit from this lesson, you must have these prerequisite skills and knowledge:

- Passed the CCIE written qualification exam
- Basic knowledge of transparent bridging and the configuration of transparent bridging

Outline

This lesson includes these sections:

- Overview
- DLSw+ Basic Configuration
- Multiple Encapsulation Options
- DLSw+ Scalability Features
- Enhanced Availability Features
- Summary
- Lesson Assessment (Quiz)

DLSw+ Basic Configuration

This section covers the basic configuration of DLSw+.

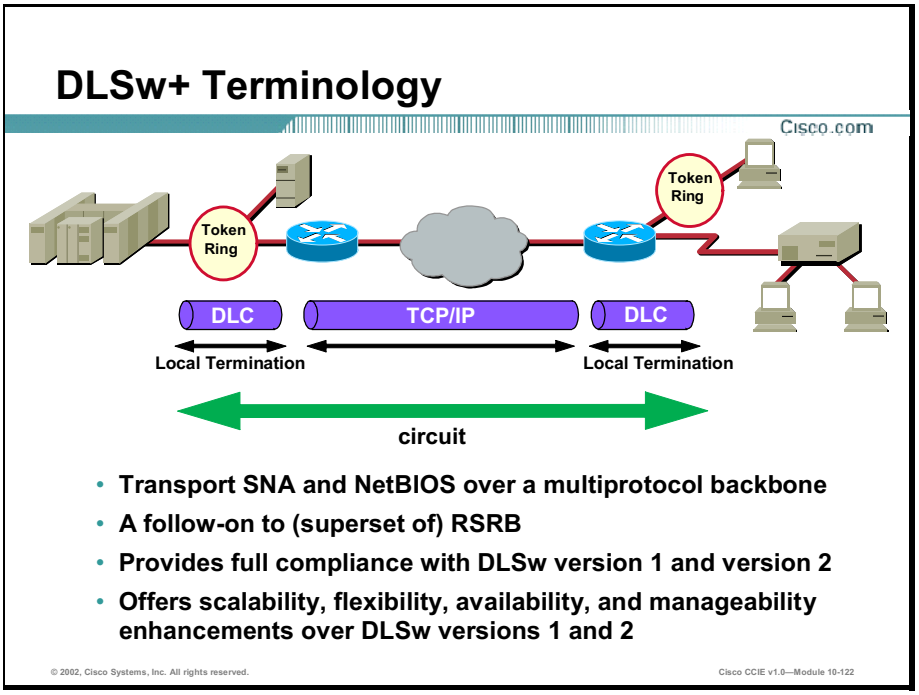

DLSw stands for Data Link Switching. DLSw is a switch-to-switch protocol that is used to transport IBM Systems Network Architecture (SNA) and IBM NetBIOS traffic over an Internet Protocol (IP) network. This protocol doesn't provide full routing, but instead provides switching at the SNA Data Link layer and encapsulation in Transmission Control Protocol/Internet Protocol (TCP/IP) for transport over the Internet.

DLSw+ is Cisco's implementation of DLSw. In addition to the DLSw standard, DLSw+ includes the following features:

- Choice of transport option, including Transmission Control Protocol (TCP), Fast Sequenced Transport (FST), and direct encapsulation.
- Scalability enhancements through: peer groups, on-demand peers, explorer firewalls, and location learning.
- Media conversion between local and remote Local Area Networks (LANs) and Synchronous Data Link Control (SDLC) or Ethernet.

DLSw+ Basic Configuration

This section details the basic configuration of a DLSw+ peer connection.

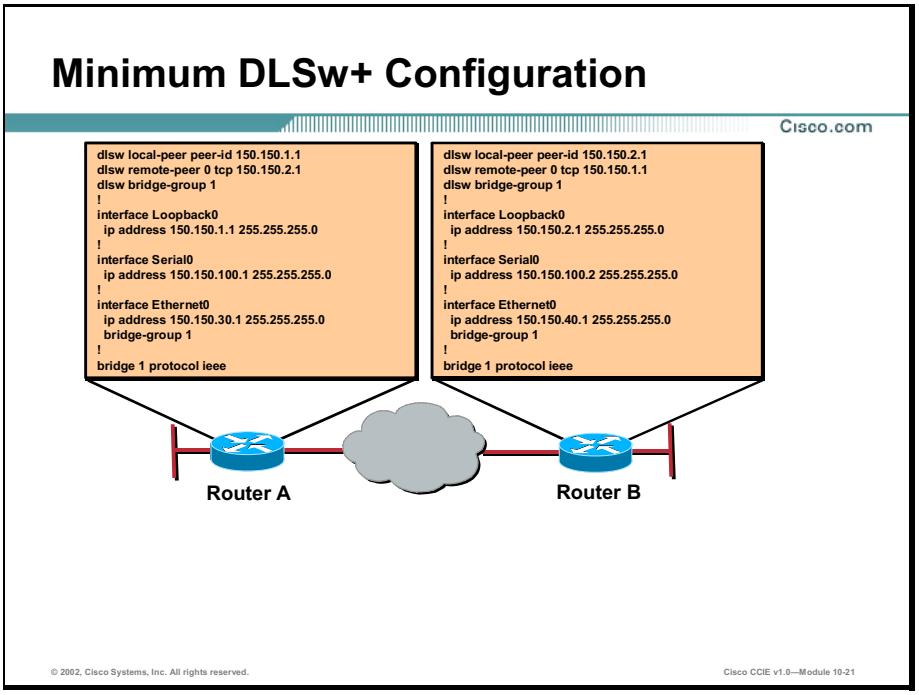

You must configure the following when configuring DLSw+ for Ethernet networks:

- Local peer (**dlsw local-peer peer-id x.x.x.x**)
- Remote peer (**dlsw remote-peer 0 tcp x.x.x.x**)
- Enable Transparent Bridging (**bridge <1-255> protocol ieee | dec**)
- Specify which Ethernet networks should be bridged (**bridge-group <1-255>**)
- Specify which bridge-groups will be carried across the DLSw+ connection (**dlsw bridge-group <1-255>**)

Note: It is recommended that you configure a loopback interface and use its IP address as the identifier for your DLSw+ local peer ID for redundancy purposes.

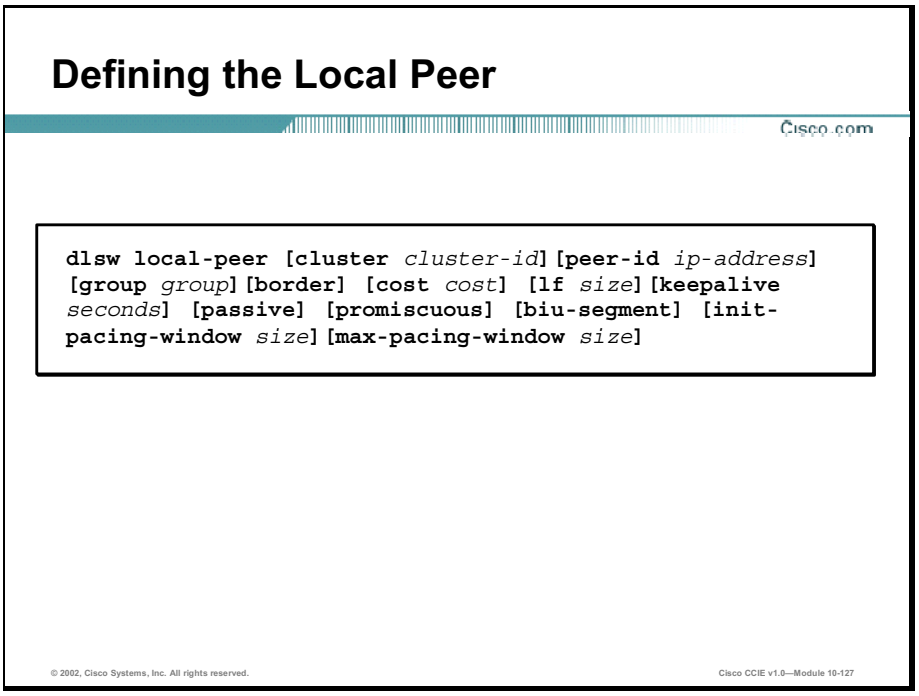

To define the DLSw+ local peer, use the global configuration command shown.

Table 9-1: dlsw local-peer Variables

Variable	Description
`cluster cluster-id`	Implements the DLSw+ Peer Clusters feature and defines the router as part of a particular cluster. The valid range is 1 to 255.
`peer-id ip-address`	Local peer IP address. This address is required when Fast-Sequenced Transport (FST) or TCP is used.
`group group`	Peer group number for this router. The valid range is 1 to 255.
`border`	Enables the router as a border peer. The group option must be specified to use the border peer option.
`cost cost`	Peer cost advertised to remote peers in the capabilities exchange. The valid range is 1 to 5.
`lf size`	Largest frame size for this local peer. Valid sizes are 516, 1470, 1500, 2052, 4472, 8144, 11407, 11454, and 17800 bytes.
`keepalive seconds`	Default remote peer keepalive interval in seconds. The valid range is 0 to 1200 seconds. The default is 30 seconds. The value 0 means no keepalives.
`passive`	Specifies that this router does not initiate remote peer connections to configured peers.
`promiscuous`	Accept connections from non-configured remote peers.
`biu-segment`	dlsw+ spoofs the maximum receivable I-frame size in XID so that each end station sends its largest frame.
`init-pacing-window size`	Size of the initial pacing window as defined in RFC 1795. The valid range is 1 to 2000.
`max-pacing-window size`	Maximum size of the pacing window as defined in RFC 1795. The valid range is 1 to 2000.

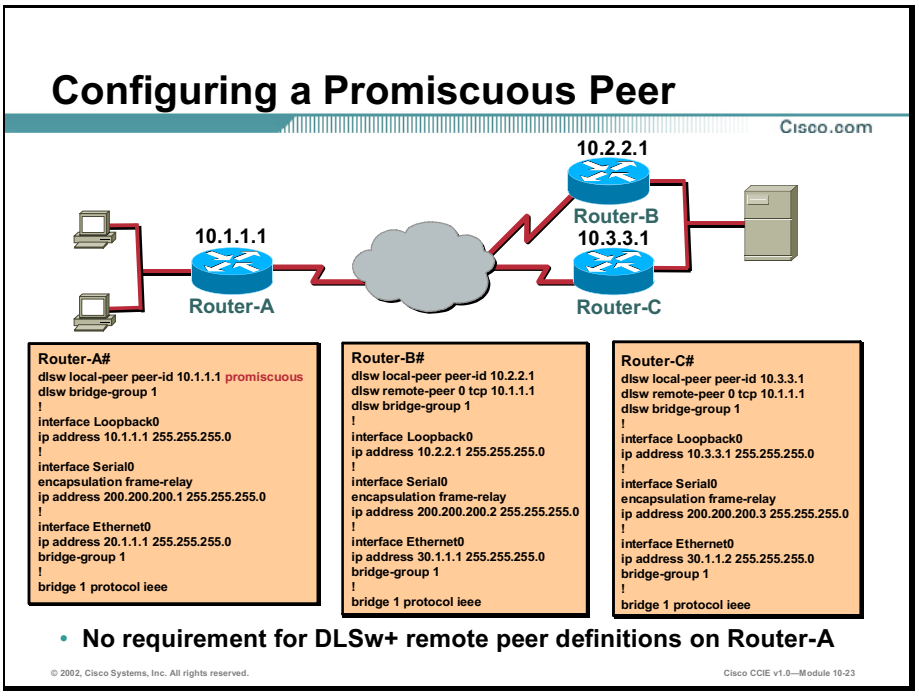

Promiscuous mode allows remote peers to connect without specifying each and every remote peer. This mode is typically used at a central site.

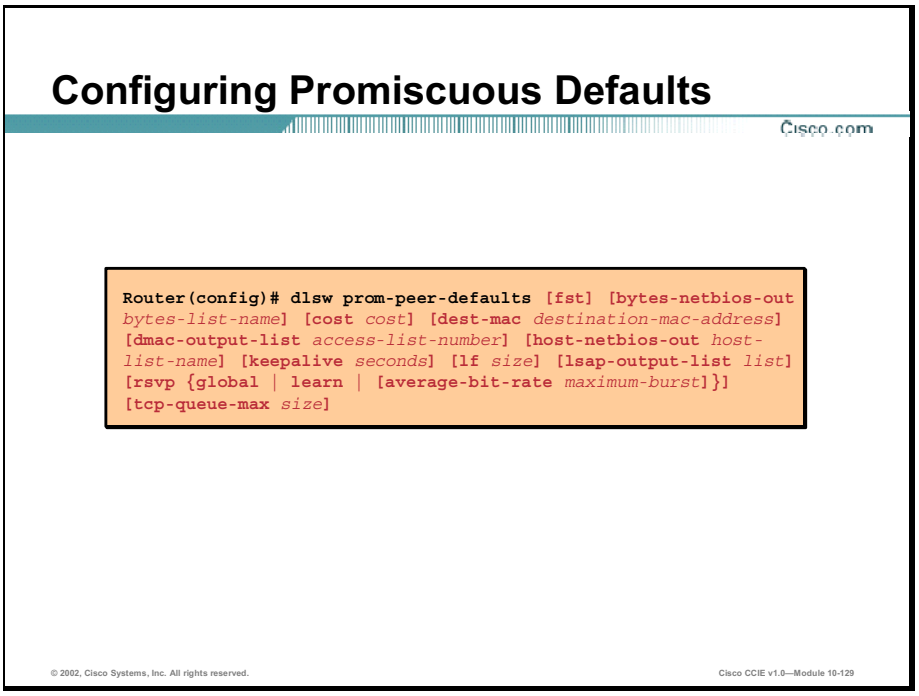

To configure promiscuous peer defaults, perform the task in global configuration mode as shown. Promiscuous mode provides for configuring remote peer options without configuring every remote peer with a 'remote peer' configuration statement.

Table 9-2: dlsw prom-peer-defaults Variables

Variable	Description
`fst`	Use FST encapsulation for all promiscuous peers established by this router.
`bytes-netbios-out` *`bytes-list-name`*	Configures NetBIOS bytes output filtering for promiscuous peers. The bytes-list-name is the name of the previously defined NetBIOS bytes access list filter.
`cost` *`cost`*	Specifies the cost to reach promiscuous peers. The valid range is 1 to 5. The default cost is 3.
`dest-mac` *`destination-mac-address`*	Specifies the exclusive destination MAC address for promiscuous peers.
`dmac-output-list` *`access-list-number`*	Specifies the filter output destination MAC addresses.
`host-netbios-out` *`host-list-name`*	Configures NetBIOS host output filtering for promiscuous peers. The host-list-name is the name of the previously defined NetBIOS host access list filter.
`keepalive` *`seconds`*	Configures the promiscuous keepalive interval. The valid range is 0 to 1200 seconds. The default is 30 seconds.
`lf` *`size`*	Largest frame size for this promiscuous peer. Valid sizes are 516, 1470, 1500, 2052, 4472, 8144, 11407, 11454, and 17800 bytes.
`lsap-output-list` *`list`*	Configures Link Service Advertising Protocol (LSAP) output filtering for promiscuous peers. Valid numbers are 200 to 299.
`rsvp global`	Sets the Resource Reservation Protocol (RSVP) parameters to the global values.

Variable	Description
`rsvp learn`	Configures RSVP parameters (average-bit-rate and maximum burst rate) to be those of the remote peer to which the promiscuous peer is connecting.
`average-bit-rate`	Configures RSVP parameters for this peer connection, which are different from the global values. Average bit rate (kilobits per second) to reserve up to 75 percent of the total bits on the interface. The valid range is 0 to 4294967.
`maximum-burst`	Maximum burst size (kilobytes of data in queue). The valid range is 0 to 4294967.
`tcp-queue-max` *`size`*	Configures the maximum output TCP queue size for promiscuous peers.

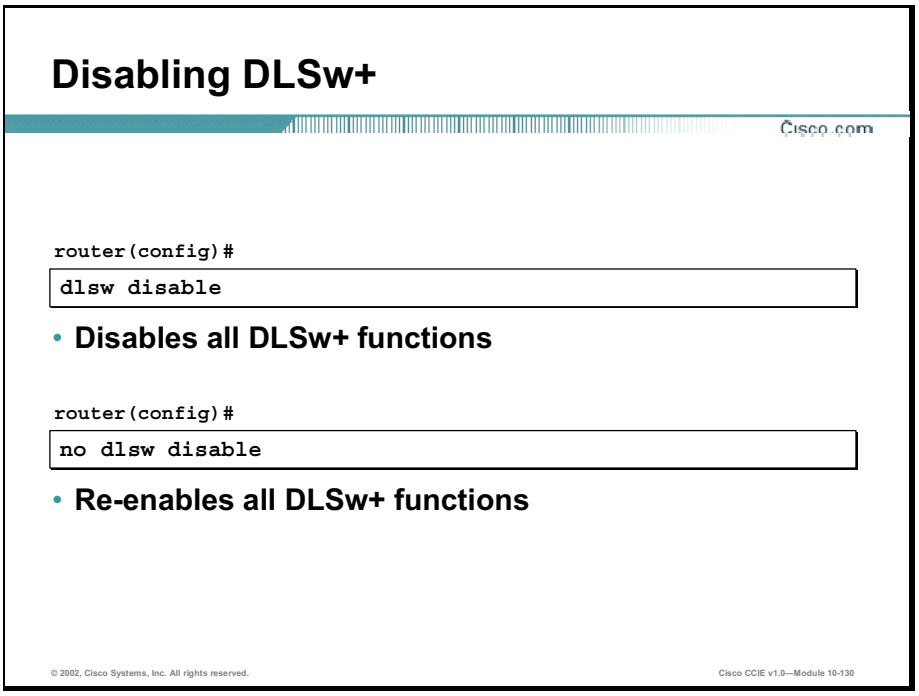

To disable all DLSw+ functions on a router without altering or removing configuration entries, use the global configuration command shown. To re-enable all DLSw+ functions, you can use the global configuration command shown.

Table 9-3: DLSw Commands

Command	Description
`dlsw disable`	Disables all DLSw+ functions
`no dlsw disable`	Re-enables all DLSw+ functions

Multiple Encapsulation Options

This section covers the various encapsulation options available.

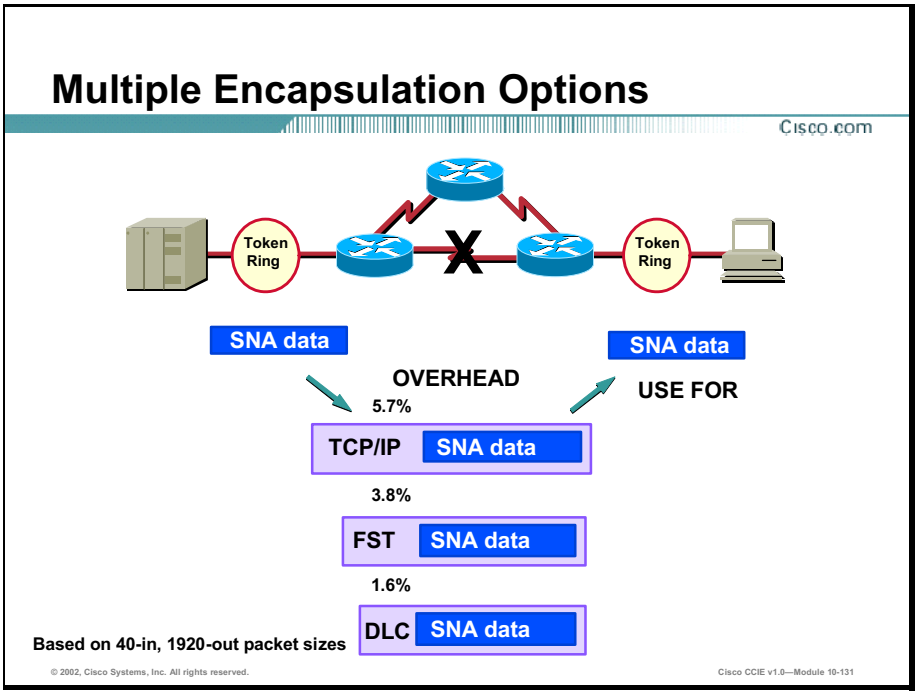

DLSw+ supports multiple encapsulation options:

- TCP/IP
- FST/IP
- Direct
- DLSw Lite

FST/IP encapsulation is fast to reroute around link failures, but does not provide local acknowledgment and resequencing.

Direct encapsulation in the appropriate link protocol minimizes overhead and is the best solution for transport across point-to-point lines (where no alternate routes are available).

When transmitting over Frame Relay, direct encapsulation uses the format described in Request for Comments (RFC) 1490, and supports optional local acknowledgment of data. Local acknowledgment not only prevents data link time outs, but also eliminates unnecessary traffic (keepalives and acknowledgments) from the Wide Area Network (WAN).

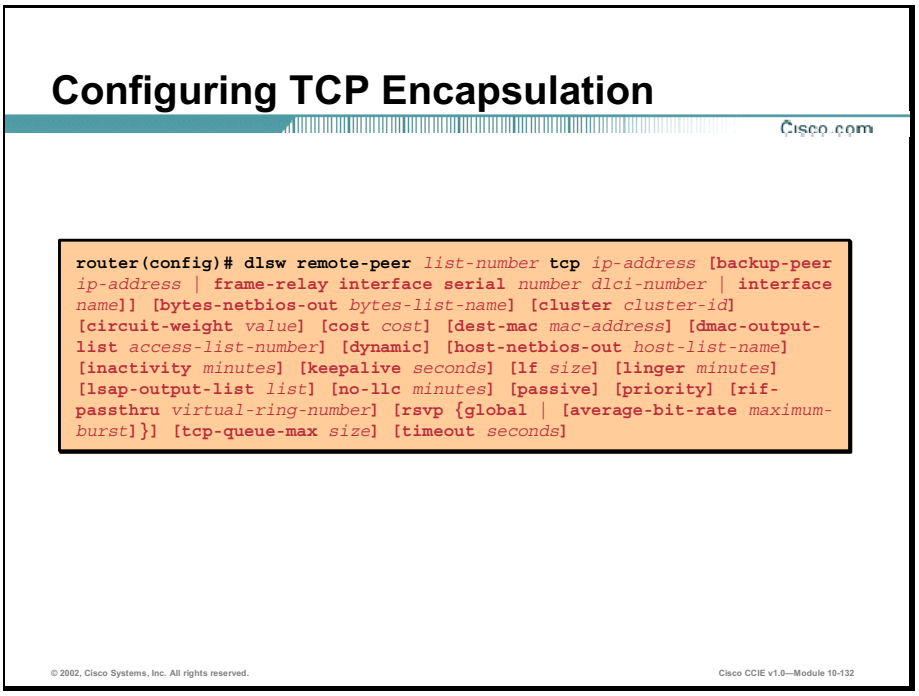

To configure TCP encapsulation for the remote peer, perform the task shown in global configuration mode.

Table 9-4: DLSw remote-peer Variables

Command	Description
`list-number`	Remote peer ring group list number. This ring group list number default is 0. Otherwise, this value must match the number you specify with the DLSw ring-list, DLSw port-list, or DLSw bgroup-list command.
`ip-address`	IP address of the remote peer with which the router is to communicate.
`backup-peer ip-address`	IP address of the existing TCP/FST peer for which this peer is the backup peer.
`backup-peer frame-relay interface serial number dlci-number`	Serial interface and Data-Link Connection Identifiers (DLCI) number of the existing Direct/LLC2 Frame Relay peer for which this peer is the backup peer.
`backup-peer interface name`	Interface name of the existing direct peer for which this peer is the backup peer.
`backup-peer circuit-inactivity minutes`	Configures the length of time a circuit is idle before terminating the circuit. The valid range is 1 to 1440.
`bytes-netbios-out` *`bytes-list-name`*	Configures NetBIOS bytes output filtering for this peer. The bytes-list-name argument is the name of the previously defined NetBIOS bytes access list filter.
`cluster` *`cluster-id`*	Used to indicate to a border peer that a particular remote-peer should be treated as part of a specific peer cluster. The valid Range is 1 to 255.
`circuit-weight` *`value`*	Configures the target state that DLSw+ tries to maintain. The valid range is 1 to 100.

`cost` *`cost`*	Cost to reach this remote peer. The valid range is 1 to 5.
`dest-mac` *`mac-address`*	Specifies the exclusive 48-bit destination MAC address, written as a dotted triple of four-digit hexadecimal numbers, for peer-on-demand peers. If the dynamic keyword is also specified, the TCP connection is established only when there is an explorer frame destined for the specified MAC address.
`dmac-output-list` *`access-list-number`*	Specifies the filter output destination MAC addresses. The access-list-number is the list number specified in an access-list command. If the dynamic keyword is also specified, the TCP connection is established only when the explorer frame passes the specified access list.
`dynamic`	Establishes the TCP connection only when there is DLSw+ data to send.
`host-netbios-out` *`host-list-name`*	Configures NetBIOS host output filtering for this peer. The host-list-name is the name of the previously defined NetBIOS host access list filter.
`inactivity` *`minutes`*	Configures the length of time a connection is idle before closing the dynamic remote peer connection. The valid range is 1 to 300 minutes. The default is 5 minutes.
`keepalive` *`seconds`*	Sets the keepalive interval for this remote peer. The range is 0 to 1200 seconds.
`lf` *`size`*	Largest frame size, in bytes, this local peer uses on a circuit to avoid segmented frames. Valid sizes are 516, 1470, 1500, 2052, 4472, 8144, 11407, 11454, and 17800 bytes.
`linger` *`minutes`*	Configures length of time the backup peer remains connected after the primary peer connection is reestablished. The valid range is 1 to 300 minutes. The default is 5 minutes.
`lsap-output-list` *`list`*	Filters output Institute of Electrical and Electronics Engineers (IEEE) 802.5 encapsulated packets. Valid access list numbers are in the range 200 to 299.
`no-llc` *`minutes`*	Configures the length of time a remote peer remains connected after all LLC2 connections are gone. The valid range is 1 to 300 minutes. The default is 5 minutes.
`passive`	Designates this remote peer as passive.
`priority`	Enables prioritization features for this remote peer. Valid TCP port numbers are the following: High—2065, Medium—1981, Normal—1982, Low—1983.
`rif-passthru` *`virtual-ring-number`*	Configures the remote peer as Routing Information Field (RIF)-Passthru. The virtual-ring-number value is the same number as the ring number value assigned in the source-bridge ring-group commands of the DLSw+ Passthru peers.
`rsvp global`	Configures the RSVP parameters for this specific peer back to the global values.
`rsvp average-bit-rate`	Configures RSVP parameters for this peer, which are different from the global values. Average bit rate (kilobits per second) reserves up to 75 percent of the total bits on the interface. Range is 0 to 4294967.
`maximum` *`burst`*	Maximum burst size (kilobytes of data in queue). Range is 0 to 4294967.
`tcp-queue-max` *`size`*	Maximum output TCP queue size for this remote peer. The valid maximum TCP queue size is a number in the range 10 to 2000.
`timeout` *`seconds`*	Resend time limit for TCP. The valid range is 5 to 1200 seconds. The default is 90 seconds.

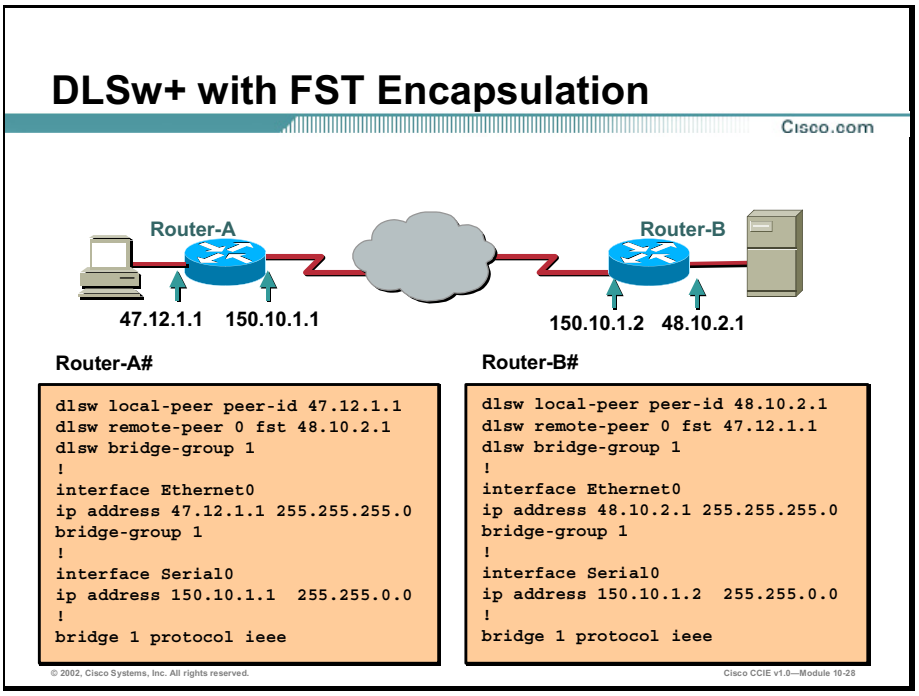

This example illustrates a DLSw+ Fast-Sequenced Transport (FST) peer configuration – notice that the configuration is very similar to a TCP peer configuration. FST encapsulation is typically used for transport across WANs with an arbitrary topology that has sufficient bandwidth to accommodate SNA and NetBIOS.

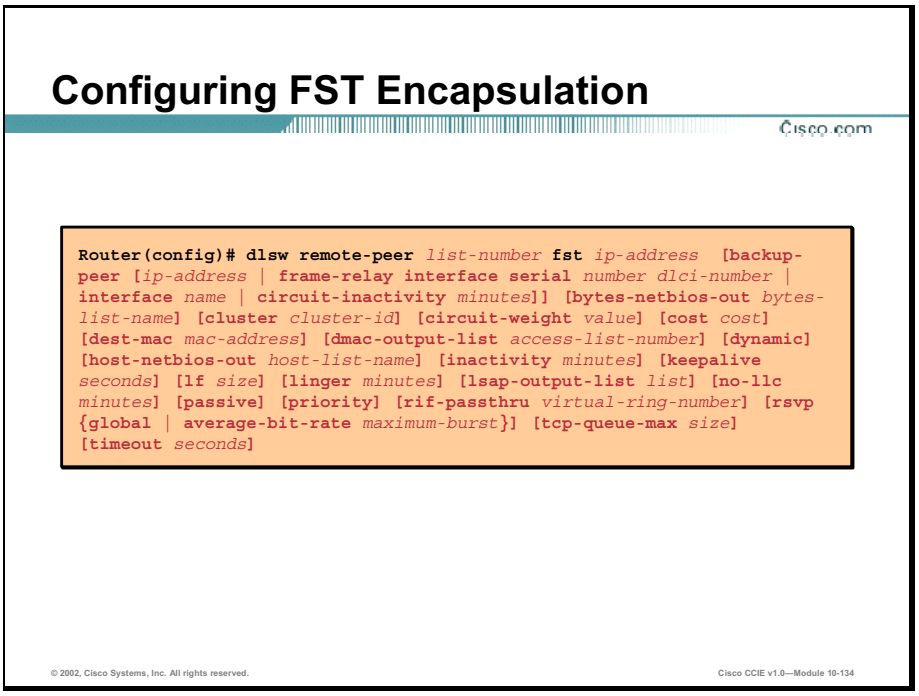

To configure FST encapsulation for the remote peer, perform the task shown in global configuration mode.

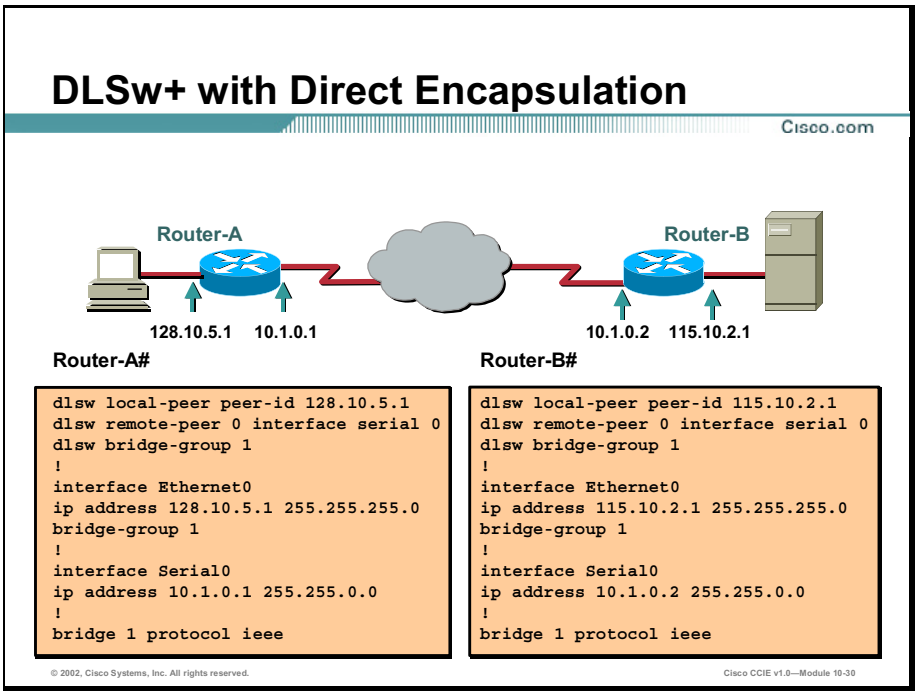

Direct encapsulation is typically used for transport across a point-to-point connection where the benefits of an arbitrary topology are not important. This example illustrates a network using direct encapsulation to transport frames from a client on the 128.10.5.0 network to a server on the 115.10.2.0 network.

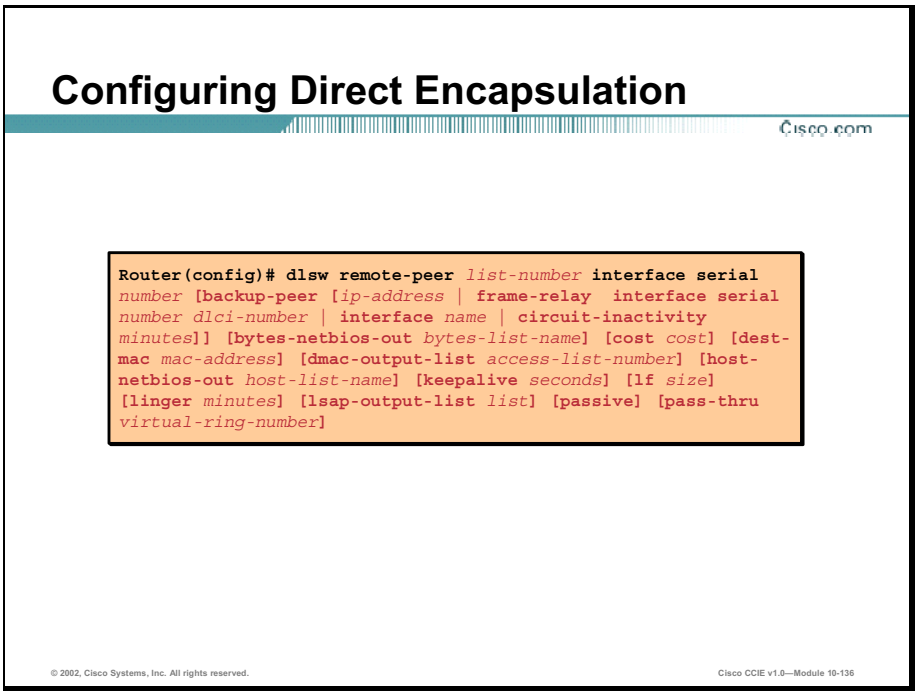

To configure direct encapsulation of either Frame Relay or HDLC, perform the task shown in global configuration mode.

Table 9-5: DLSw remote-peer serial Variables

Command	Description
list-number	Ring list number. The valid range is 1 to 255. The default is 0, which means all.
serial number	Specifies the remote peer by direct serial interface.
backup-peer *ip-address*	IP address of the existing TCP/FST peer for which this peer is the backup peer.
backup-peer frame-relay interface serial *number dlci-number*	Serial interface and DLCI number of the existing Direct/LLC2 frame-relay peer for which this peer is the backup peer.
backup-peer interface *name*	Interface name of the existing direct peer for which this peer is the backup peer.
backup-peer circuit-inactivity *minutes*	Configures the length of time a circuit is inactive before being terminated. May be used with the linger option. The valid range is 1 to 1440 minutes.
bytes-netbios-out *bytes-list-name*	Configures NetBIOS bytes output filtering for this peer. The bytes-list-name argument is the name of the previously defined NetBIOS bytes access list filter.
cost cost	Cost to reach this remote peer. The valid range is 1 to 5.
dest-mac *mac-address*	Permits the connection to be established only when there is an explorer frame destined for the specified 48-bit MAC address written as a dotted triple of four-digit hexadecimal numbers.
dmac-output-list *access-list-number*	Permits the connection to be established only when the explorer frame passes the specified access list. The access-list-number is the list number specified in the access-list command.

Command	Description
`host-netbios-out` *host-list-name*	Configures NetBIOS host output filtering for this peer. The host-list-name is the name of the previously defined NetBIOS host access list filter.
`keepalive seconds`	Sets the keepalive interval for this remote peer. The range is 0 to 1200 seconds.
`lf size`	Largest frame size, in bytes, this local peer will use on a circuit to avoid segmented frames. Valid sizes are 516, 1470, 1500, 2052, 4472, 8144, 11407, 11454, and 17800 bytes.
`linger` *minutes*	Configures length of time the backup peer remains connected after the primary peer connection is reestablished. The valid range is 1 to 300 minutes. The default is 5 minutes.
`lsap-output-list` *list*	Filters output IEEE 802.5 encapsulated packets. Valid access list numbers are in the range 200 to 299.
`passive`	Designates this remote peer as passive.
`pass-thru`	Selects passthrough mode. The default is local acknowledgment mode.

DLSw+ Scalability Features

One of the most significant factors that limit the size of LAN internetworks is the amount of explorer traffic that traverses the WAN.

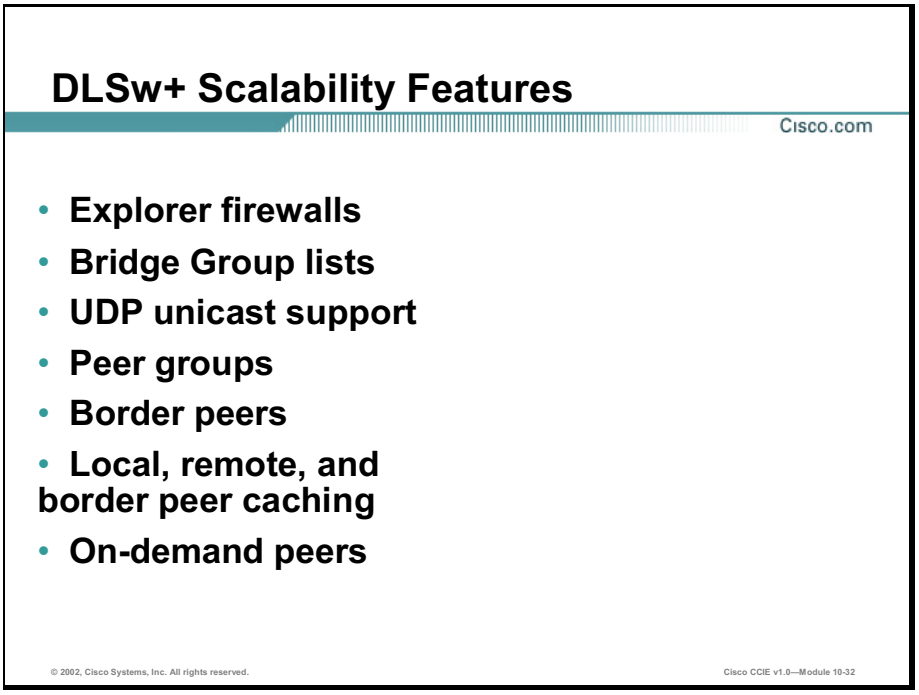

One significant factor that limits the size of SNA and NetBIOS internetworks is the amount of explorer traffic that traverses the WAN. DLSw+ includes the following features to reduce the number of explorers:

- **Peer groups:** The large SNA and NetBIOS internetworks that Cisco has helped to build over the last several years have all followed a similar structure. That structure is a hierarchical grouping of routers based upon the usual flow of broadcasts through the network. Clusters of routers in a region or a division of a company are combined into a peer group.

- **Border peers:** Within a peer group, one or more routers are designated as border peers. When a DLSw+ router receives a test frame or NetBIOS name query, it sends a single explorer frame to its border peer. The border peer takes complete responsibility for forwarding the explorer on behalf of the peer group member. This arrangement eliminates duplicate explorers on the access links and minimizes the processing required in access routers.

- **On-demand peers:** On-demand peers greatly reduce the number of peers that must be configured. It permits casual, any-to-any connection without the burden of configuring the connection in advance. It also allows any-to-any routing in large internetworks where persistent TCP connections would not otherwise be possible.

- **Explorer firewalls:** While an explorer is outstanding and waiting for a response from the destination, subsequent explorers for that MAC address are merely stored. Once the explorer response is received at the originating DLSw+, all explorers receive an immediate local response. This eliminates the start-of-day explorer storm that many networks experience.

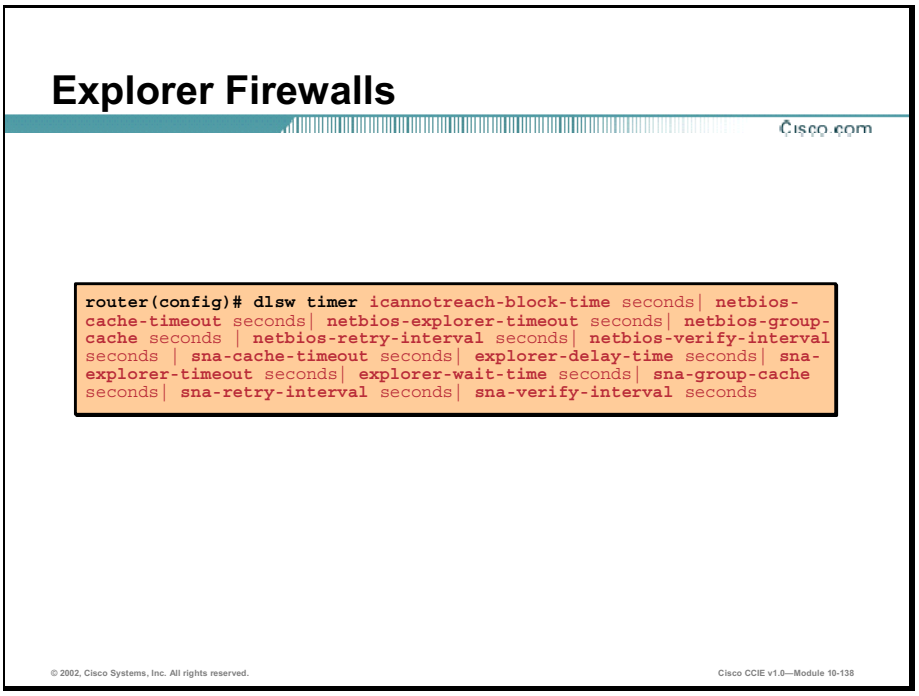

An explorer firewall controls explorers for a particular destination MAC address or NetBIOS name. This eliminates the start-of-day explorer storm that many networks experience.

Explorer timeouts can be adjusted using the following:

dlsw timer netbios-explorer-timeout *seconds*

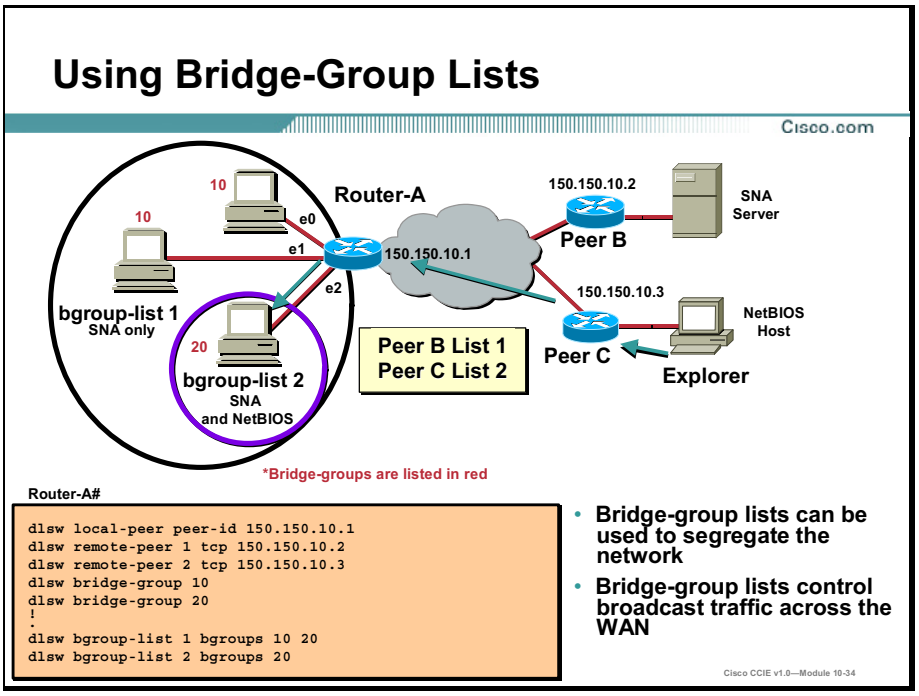

DLSw+ bridge group lists map traffic on the local Ethernet bridge group interface to remote peers. You can create a bridge group list and apply the list to remote peer definitions. Traffic received from a remote peer is only forwarded to the bridge group specified in the bridge group list. Traffic received from a local interface is only forwarded to peers if the input bridge group number appears in the bridge group list applied to the remote peer definition. The definition of a bridge group list is optional. To define a bridge-group list, use the following command in global configuration mode:

Table 9-6: DLSw Commands

Command	Description
`dlsw disable`	Disables all DLSw+ functions
`no dlsw disable`	Re-enables all DLSw+ functions

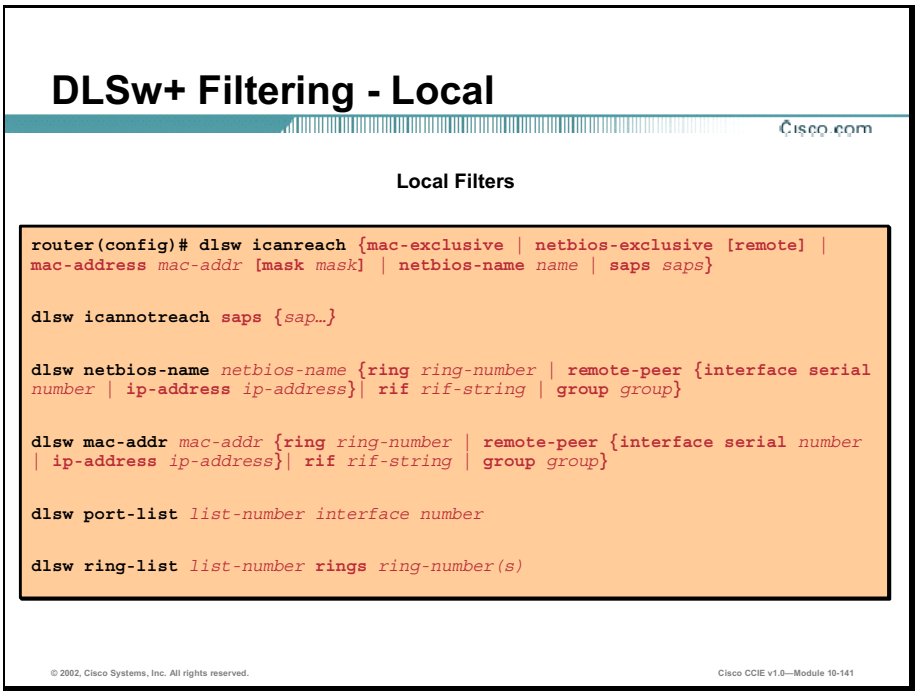

DLSw+ local filtering is accomplished by using the code shown.

Example 1, **dlsw icanreach**: This example on the central site router will provide all remote peers with an entry on their DLSw+ reachability table for the host MAC address pointing to the central router IP address. This entry is in the UNCONFIRM state, which indicates that if the remote office router receives a local test or null exchange identification (XID) for the host, it sends a **CUR_ex** (Can U Reach Explorer) message to the central router only. The mask represents an 'exact match' for the MAC address.

The option 'mac-exclusive' ensures that only packets destined for the MAC addresses defined (in this case 4000.3745.0000) are allowed at the central location.

```
hostname central_site
!
dlsw local-peer peer-id 1.1.1.1
dlsw remote-peer 0 tcp 1.1.1.2
dlsw icanreach mac-exclusive
dlsw icanreach mac-address 0001.0374A.2764 mask ffff.ffff.ffff
!
```

Example 2, **dlsw icannotreach saps**: The **dlsw icannotreach saps** command on the central site route allows you to filter those protocols you know are not allowed to be sent to the remote peers. If you know only what must be explicitly denied, use the **dlsw icannotreach saps** command on the central router(s), as shown in the configurations below. This example will prevent NetBIOS (sap F0) traffic. Additional saps could be listed separated by spaces. The **show dlsw capabilities** will indicate the 'unsupported saps'

In contrast, the '**dlsw icanreach saps** 04 08 0C' will allow only those saps and the **show dlsw capabilities** will indicate all other saps as 'unsupported'.

```
hostname central_site
!
dlsw local-peer peer-id 1.1.1.1
dlsw remote-peer 0 tcp 1.1.1.2
dlsw icannotreach sap F0
!
```

Example 3, **dlsw mac-addr**: In this example, you manually configure each remote router and direct them to the desired central router when looking for specific MAC addresses. This reduces unnecessary traffic going to the wrong peer. If the remote office only has one remote peer configured, then this configuration will not be beneficial. However, if multiple remote peers are configured, this configuration directs the remote site router to the right place without wasting WAN bandwidth.

```
hostname remote_site
!
dlsw local-peer peer-id 1.1.1.2
dlsw remote-peer 0 tcp 1.1.1.1
dlsw remote-peer 0 tcp 2.2.2.1
dlsw mac-addr 0001.0374A.2764 remote-peer ip-address 1.1.1.1
!
```

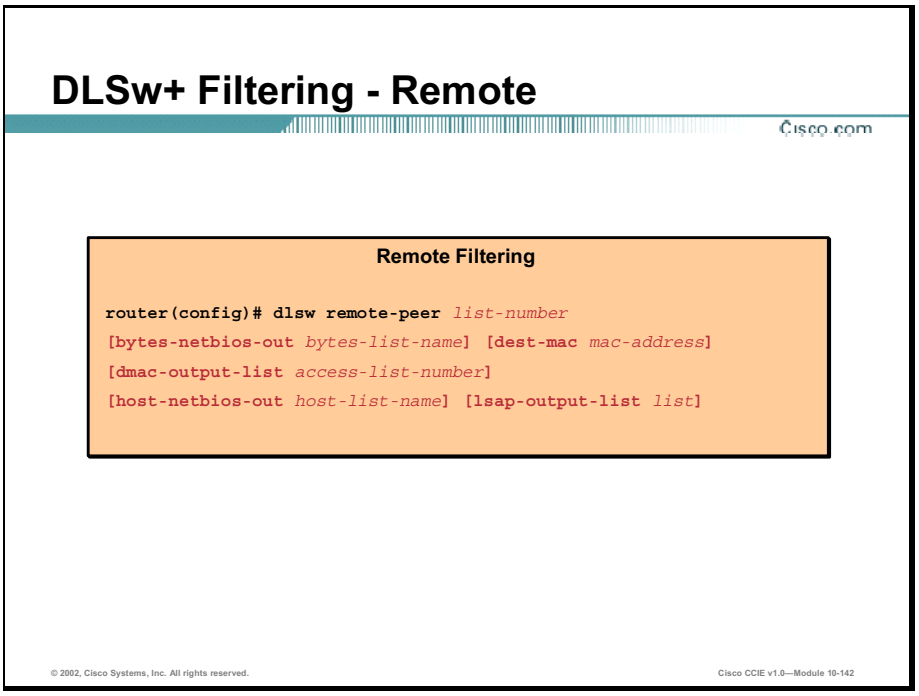

DLSw+ remote filtering is accomplished by using the options as shown.

Example 1, **lsap-output-list**: Using this method, all remote offices must be configured with the **lsap-output-list** option. No other configuration changes are required on the central router. The **lsap-output-list** links to a SAP access list (SAP ACL) that currently only allows SNA SAPs (for example, 0x00, 0x04, 0x08, and so on) to go toward the central router, and denies everything else.

The low order bit (0x01) must be on in the mask for the command / response bit. A SAP of 04 sent as a command will return as 05, indicating a response. SAP 0x00 is used in some explorers.

```
hostname remote_site
!
dlsw local-peer peer-id 1.1.1.2
dlsw remote-peer 0 tcp 1.1.1.1 lsap-output-list 200
dlsw bridge-group 1
!
interface Serial0/1
  ip address 1.1.1.2 255.255.255.0
  no ip directed-broadcast
!
access-list 200 permit 0x0000 0x0D0D
access-list 200 deny   0x0000 0xFFFF
!
```

Example 2, **dmac-output-list**: Using this method, all remote offices must be configured with the **dmac-output-list** option. No other configuration changes are required in the central router. The **dmac-output-list** links to a MAC access list that currently only allows the SNA destination MAC address(s) to go toward the central router, and denies everything else. The mask field is a wild card mask; a one bit indicates 'do not care.'

```
hostname remote_site
!
dlsw local-peer peer-id 1.1.1.2
dlsw remote-peer 0 tcp 1.1.1.1 dmac-output-list 701
dlsw bridge-group 1
!
interface Serial0/1
  ip address 1.1.1.2 255.255.255.0
  no ip directed-broadcast
!
access-list 701 permit 0001.0374A.2764 0000.0000.0000
access-list 700 deny   0000.0000.0000 FFFF.FFFF.FFFF
!
```

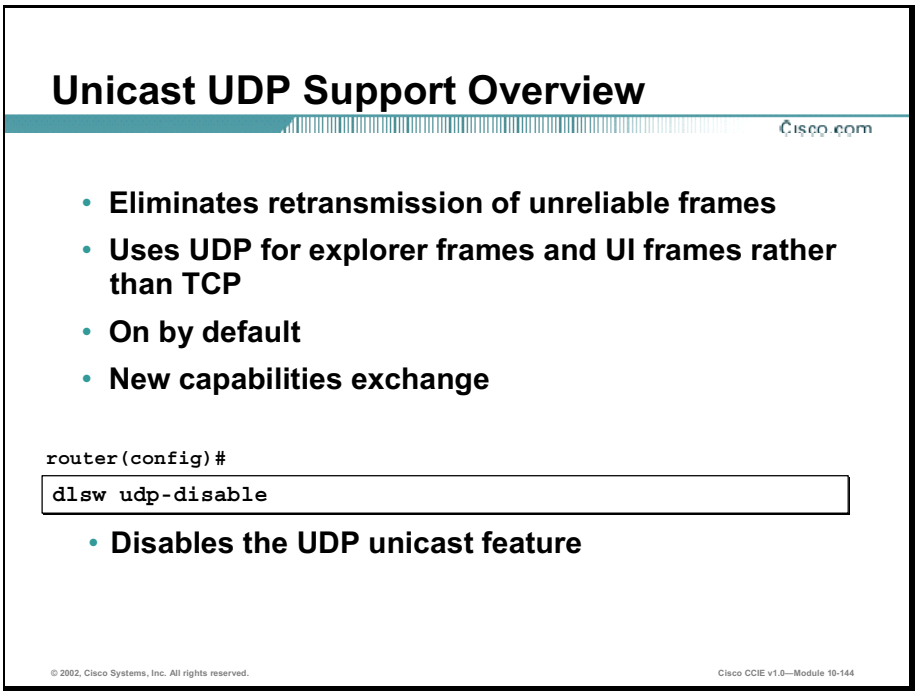

With User Datagram Protocol (UDP) Unicast, explorer frames and unnumbered information frames are sent via UDP rather than TCP. This feature eliminates retransmission of explorers and unnumbered information frames that could occur during congestion. Cisco's DLSw+ introduced the UDP Unicast feature prior to the release of DLSw Version 2 in Cisco IOS Release 11.2(6)F.

One difference between the two enhancements is that the Release 11.2(6)F UDP Unicast feature requires that a TCP connection exist before packets are sent via UDP. Because the TCP session is up and capabilities have been exchanged, the peers know exclusive reachability information that will permit them to further reduce the explorer load on the network.

DLSw+ Version 2, on the other hand, sends UDP/IP multicast and unicast before the TCP connection exists. Although DLSw+ Version 2 employs IP multicast service when address resolution packets (CANUREACH_EX, NETBIOS_NQ_ex, NETBIOS_ANQ, and DATAFRAME) are sent to multiple destinations, the response frames (ICANREACH_ex and NAME_RECOGNIZED_ex), are sent via UDP unicast.

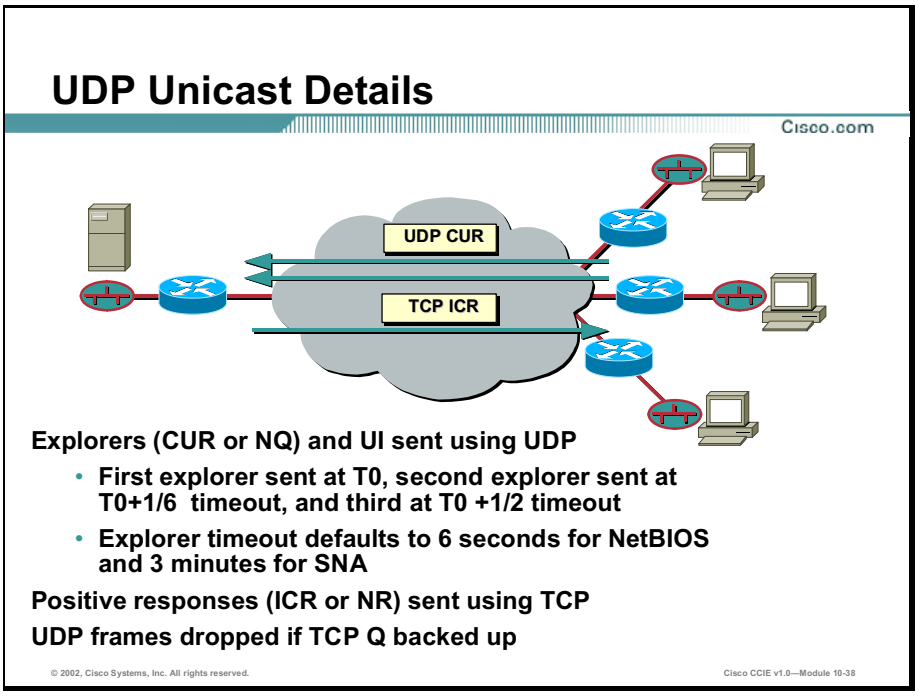

The UDP Unicast feature eliminates the retransmission of explorer frames and unnumbered information frames that occur during congestion.

UDP Unicast abbreviations:

- CUR is CANUREACH
- NQ is Name Query
- ICR is ICANREACH
- NR is Name Recognized

The Explorer Timeouts use existing DLSw timers. Their meaning does not change – they still represent the time DLSw will wait for a response before deleting the entry. However, DLSw+ is used to send only a single explorer. Now, because the delivery vehicle is unreliable, DLSw+ will send up to three explorers.

The second explorer is sent T0 + 1/6 timeout, or T0 + 1 second for NetBIOS, T0 + 30 seconds for SNA.

The third explorer is sent at T0 + 3 seconds for NetBIOS and T0 + 90 seconds for SNA.

Explorer timeouts can be adjusted using the command shown.

```
dlsw  timer netbios-explorer-timeout seconds
```

Why Use Peer Groups?

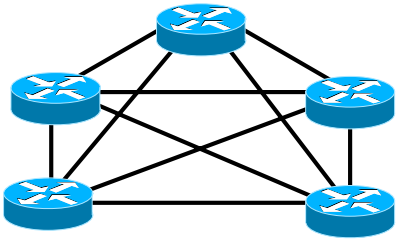

Any-to-any networks require that all routers must be peered

- **Extensive configuration**
- **Excessive explorer traffic and overhead**

Cisco CCIE v1.0—Module 10-146

A significant design limitation of RFC 1795 is shown. Version 1 required that all routers in a mesh network had to be peers. The diagram illustrates the complexity of building an any-to-any connectivity using the DLSw V1 standard. The complexity arises from the requirement for TCP connections between every pair of routers that may ever need to communicate.

These TCP connections are used not only to transport data, but also to find end systems. In a network requiring fully meshed connectivity, this mechanism does not scale well because most routers can only handle 100 to 200 TCP connections.

Not only are a large number of persistent TCP connections required, but each explorer must be replicated once per connection. This places additional processing requirements on routers and additional traffic on access links.

Finally, the configuration process required for any-to-any peer connections is extensive.

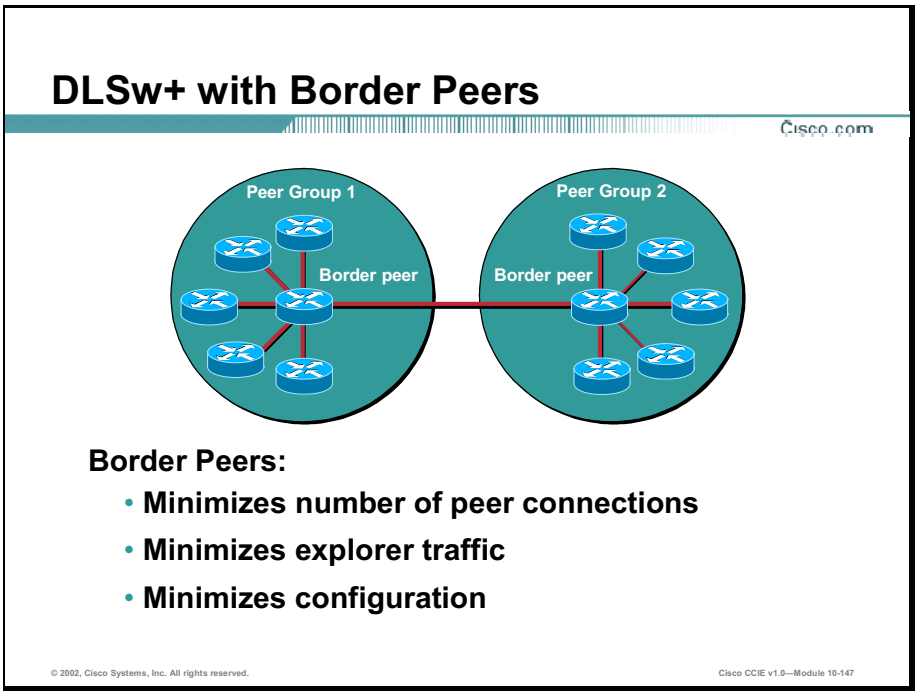

DLSw+ with Border Peers minimizes number of peer connections, explorer traffic, and configuration.

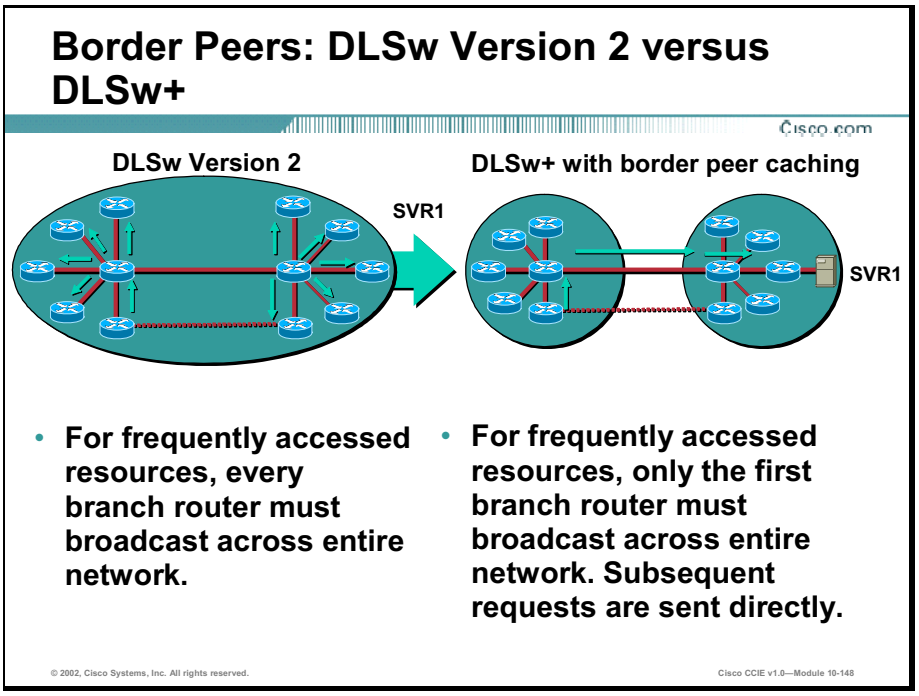

The left side of this diagram illustrates the complexity of building an any-to-any network using the DLSw standard. The complexity arises from the requirement for TCP connections between every communicating pair of routers. The lines connecting each pair of routers represent the required TCP connections for any-to-any communication.

DLSw+ allows you to group routers together and use a centralized router to do all searches. This minimizes the required configurations, the number of required TCP connections, and the broadcast traffic.

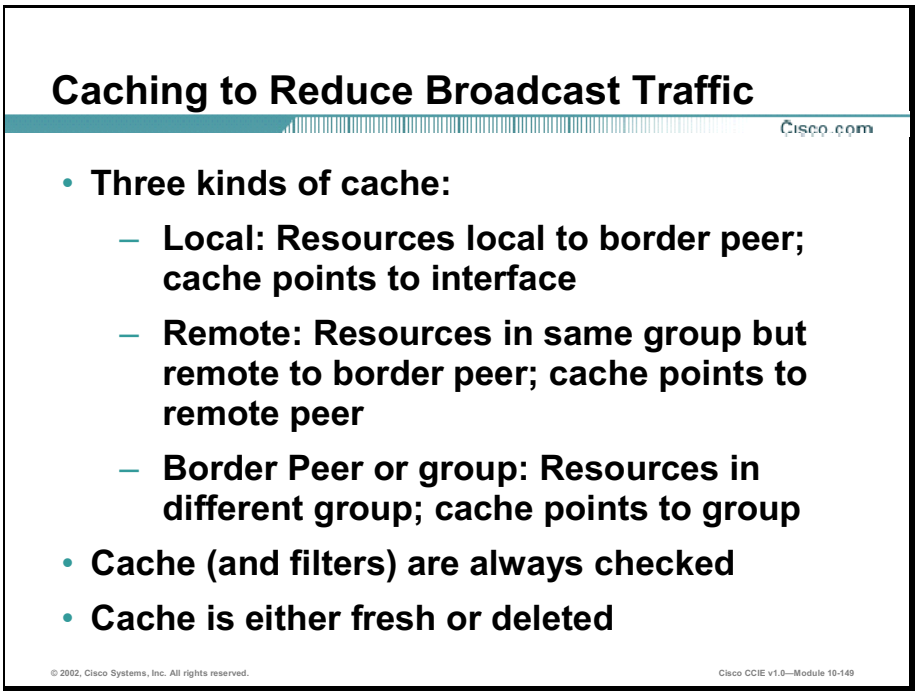

Version 2 of the DLSw standard attempts to address the problem that DLSw+ addressed with peer groups and border peers – namely, explorer replication in a meshed network. The Advanced Peer-to-Peer Networking (APPN) Implementers' Workshop (AIW) chose to use IP multicasts to find intermediate IP routers for explorer forwarding instead of DLSw-aware border peers. With the exception of using UDP for User Interface (UI) frames, the standard does not provide any additional features beyond what DLSw+ does today. In addition, by using multicasts instead of border peers, the standard complicates the task of centralizing caching, a key feature in reducing broadcast traffic.

Border Peer Caching takes advantage of the fact that the routers responsible for explorer-forwarding have DLSw capability and cache information. Border peers can either be pre-loaded or can learn the location of commonly requested resources, and use this information to eliminate subsequent broadcast traffic for those resources.

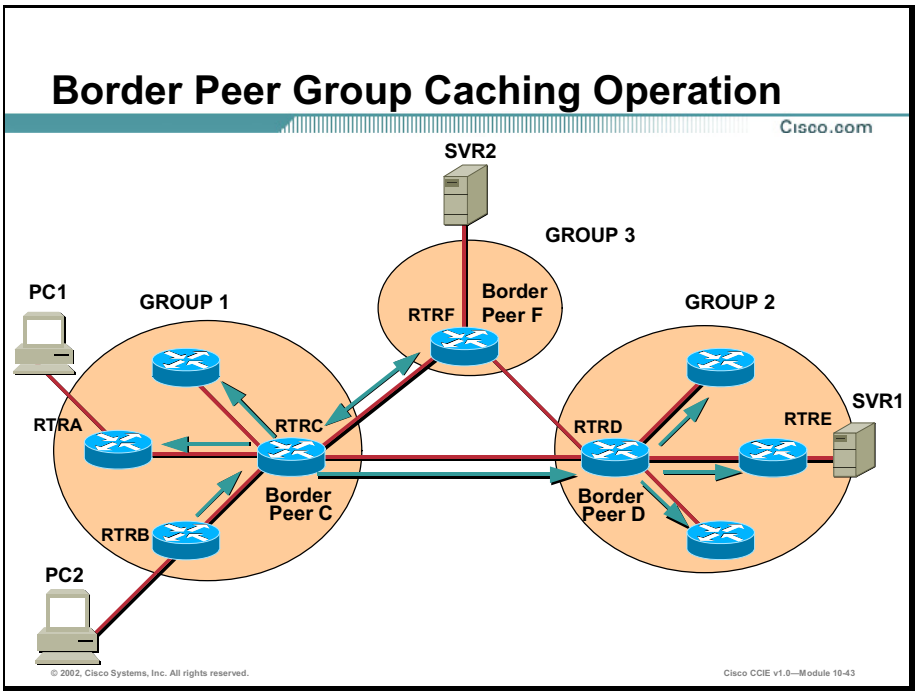

Explorers search for the destination address. If no cached address exists, the explorer is passed to border peers and on to adjacent groups.

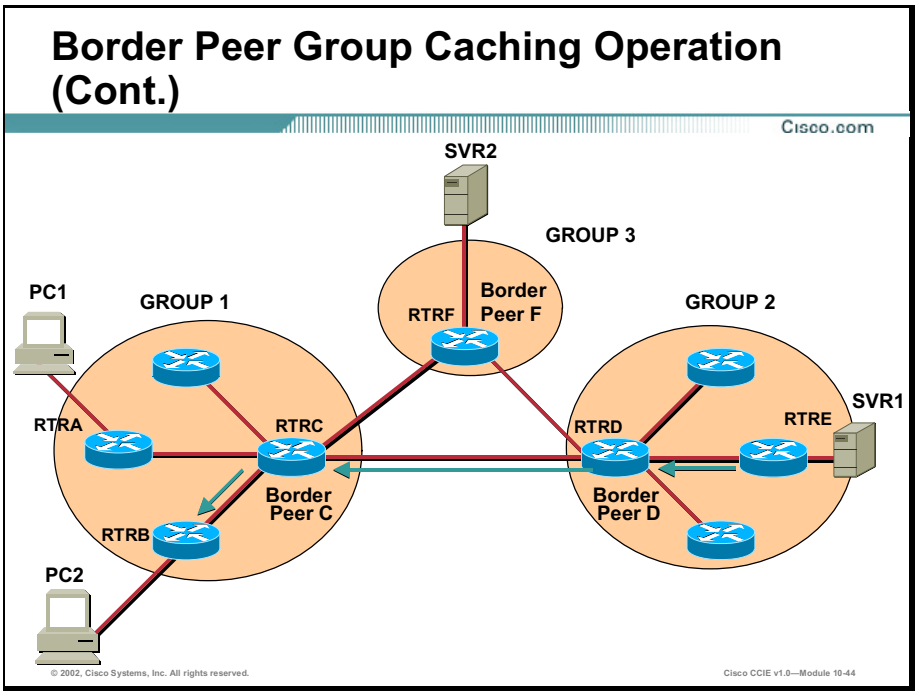

As a positive response to an explorer traversing the network, however, each border peer builds a cache entry for that resource. Only the destination resource is cached (you do not want to cache all NetBIOS clients, only the servers). Same group resources are cached in either the local or remote cache. Different group resources are cached in the group cache.

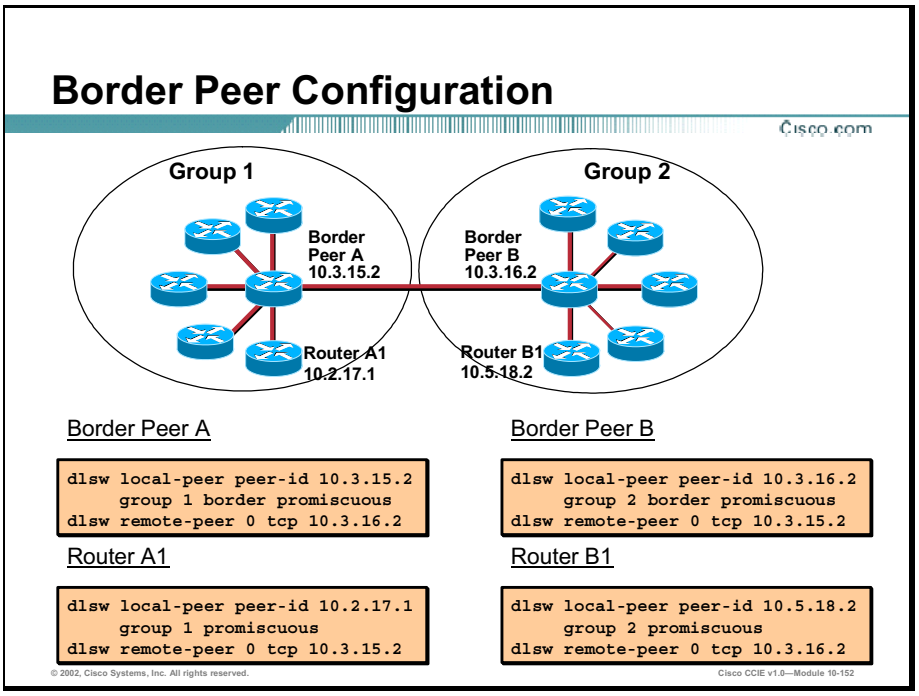

The configuration of DLSw+ border peers and peer groups is shown.

Use the **group** keyword to define a peer group, the **border** keyword to define a border peer, and the **cluster** keyword to assign the local peer to a peer cluster. When the user defines the **cluster** option in the **dlsw local-peer** command on the member peer router, the cluster information is exchanged with the border peer during the capabilities exchange as the peers become active. The border peer uses this information to make explorer replication and forwarding decisions.

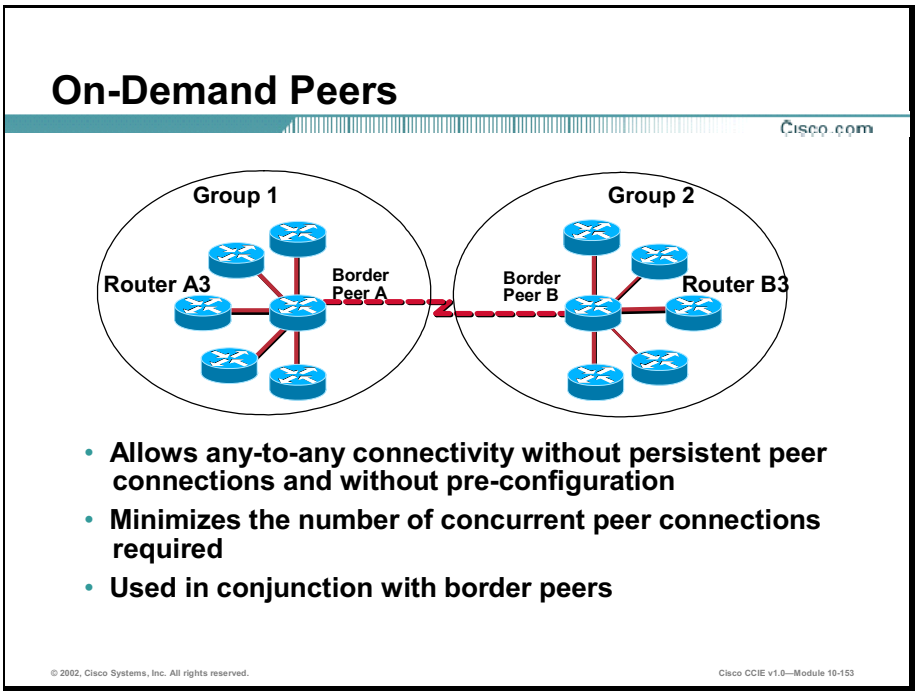

With border peers in place, it is possible for two peers to communicate with each other even though neither has a configuration for the other. This is because they learn about each other via their respective border peers.

For that reason, there is a statement that defines how to connect with peers when no DLSw remote-peer commands are used. That statement is the **dlsw peer-on-demand-defaults tcp** command, and it can be used to specify the encapsulation, filters, and timers. On-demand peer connections can be made between two peers in the same group or two peers in different groups.

Enhanced Availability Features

This section will cover enhanced peer availability features.

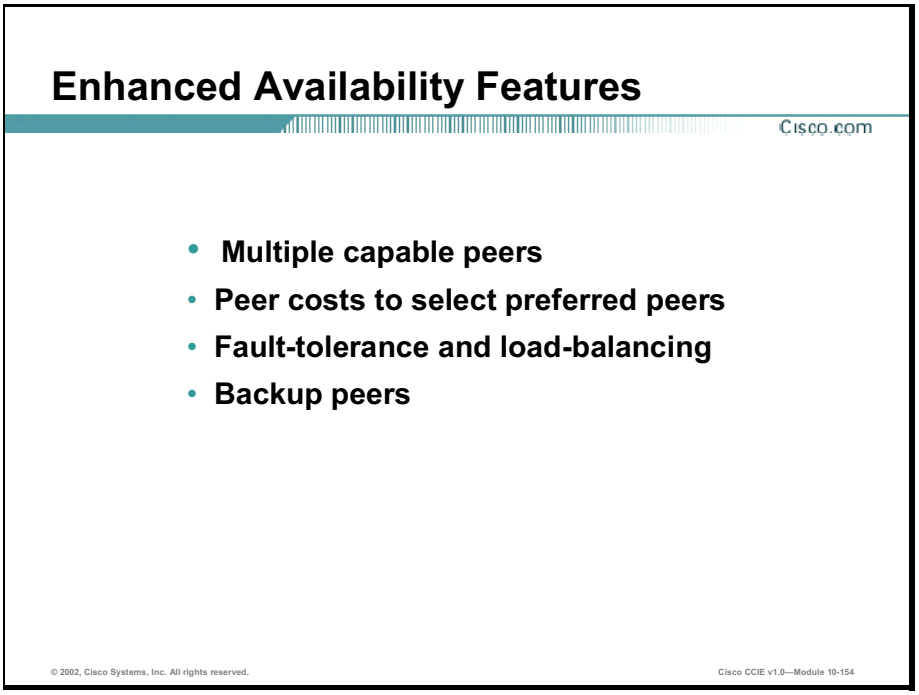

One of the ways DLSw+ offers enhanced availability is by maintaining a reachability cache of multiple paths for local and remote destination MAC addresses or NetBIOS names. For remote resources, the path specifies the peer to use to reach this resource. For local resources, the path specifies a port number. If there are multiple paths to reach a resource, the router will mark one path preferred and all other paths capable. If the preferred path is not available, the next available path is promoted to the new preferred path, and recovery over an alternative path is initiated immediately. The way that multiple capable paths are handled with DLSw+ can be biased to meet the needs of the network:

- **Fault tolerance**: Biases circuit establishment over a preferred path, but also rapidly reconnects on an active alternative path if the preferred path is lost
- **Load balancing**: Distributes circuit establishment over multiple DLSw+ peers in the network or ports on the router

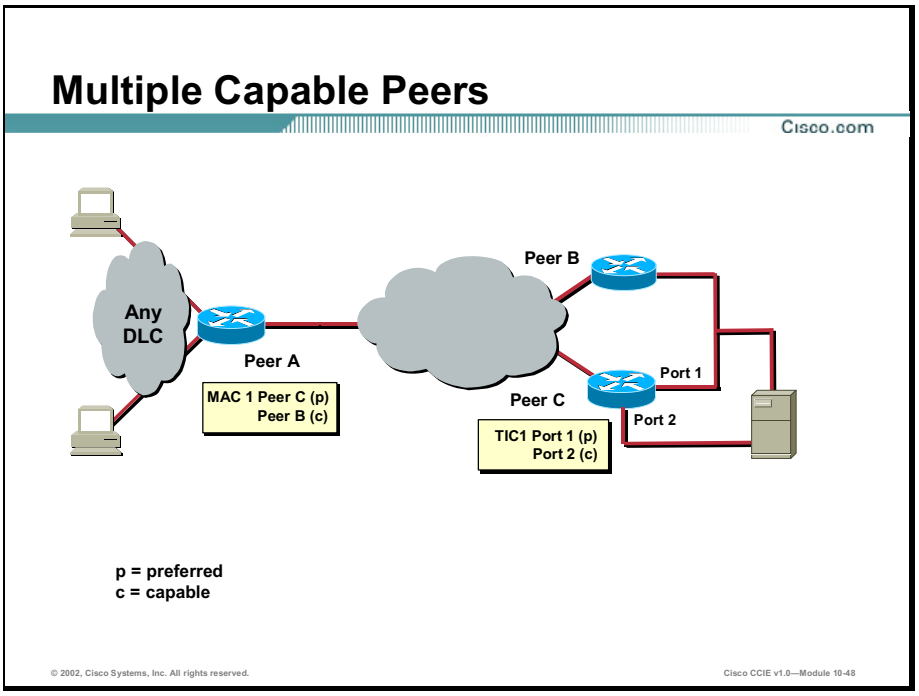

One of the enhanced availability features is a multiple capable peer. Each Cisco router maintains a preferred route and one or more capable routes to DLSw+ border peers. Should the preferred route fail, the next available route is promoted to the preferred route. No broadcasts are required, and recovery is immediate. You can use multiple capable peers for fault tolerance or for load balancing.

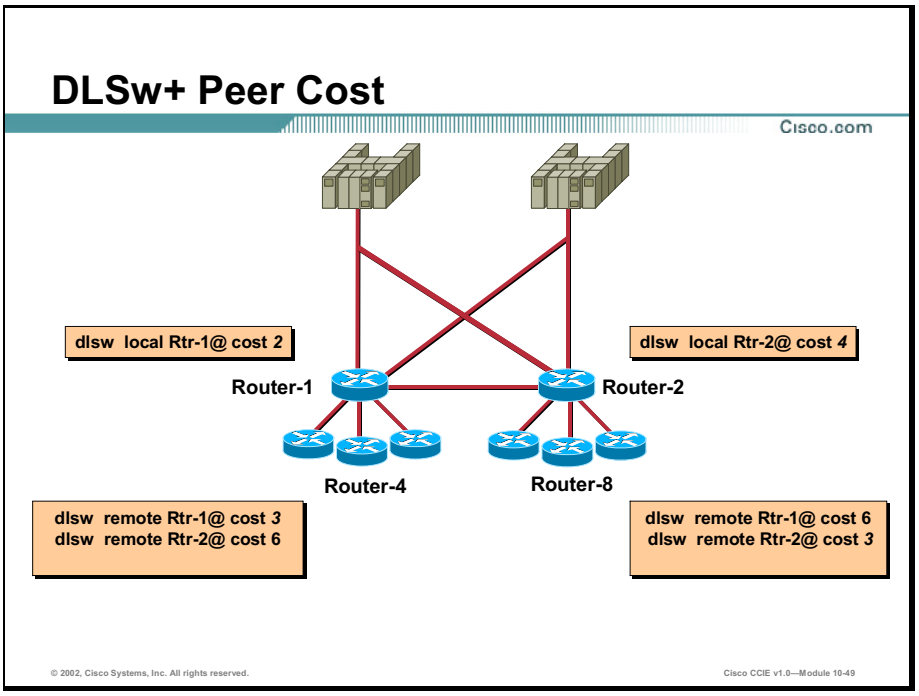

Peer costs are used to determine which router is the preferred router in the reachability cache. If no peer costs are specified, the first router to positively respond to an explorer is the preferred router for that destination.

In this example, Routers 1 and 2 specify peer costs of 2 and 4 respectively in their local peer statement. These peer costs are shared with any connecting peer as part of the capabilities exchange. If a remote router in this example has no remote peer statement for routers 1 or 2, it would always select RTR1 as the preferred peer because Router-1 has a lower cost.

RTR4 specifies a remote peer cost of 3 for Router-1 and 6 for Router-2. Therefore Router-4 always prefers Router-1 over Router-2.

Router-8, on the other hand, specifies a remote peer cost of 6 for Router-1 and 3 for Router-2. Therefore Router-8 always prefers Router-2 over Router-1.

Note	The peer cost specified on the remote peer statement overrides the peer cost specified in Routers 1 and 2 and learned in the capabilities exchange. This allows different preferred routers to be specified based on geographical region or any other logical or physical delineation.

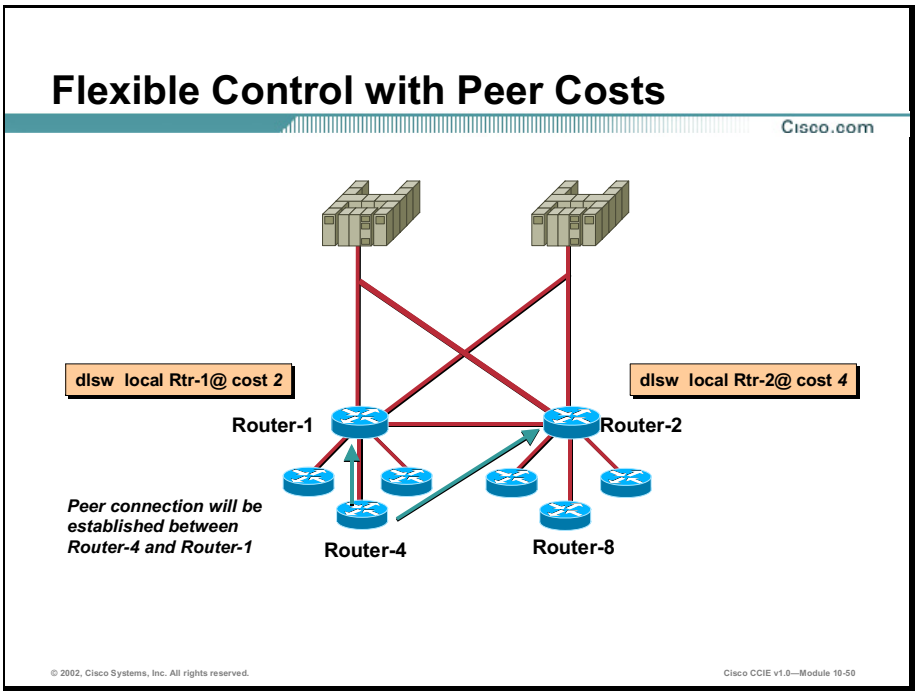

If cost is configured on the DLSw+ local peer and not on the DLSw+ remote peer, then local peer cost will apply.

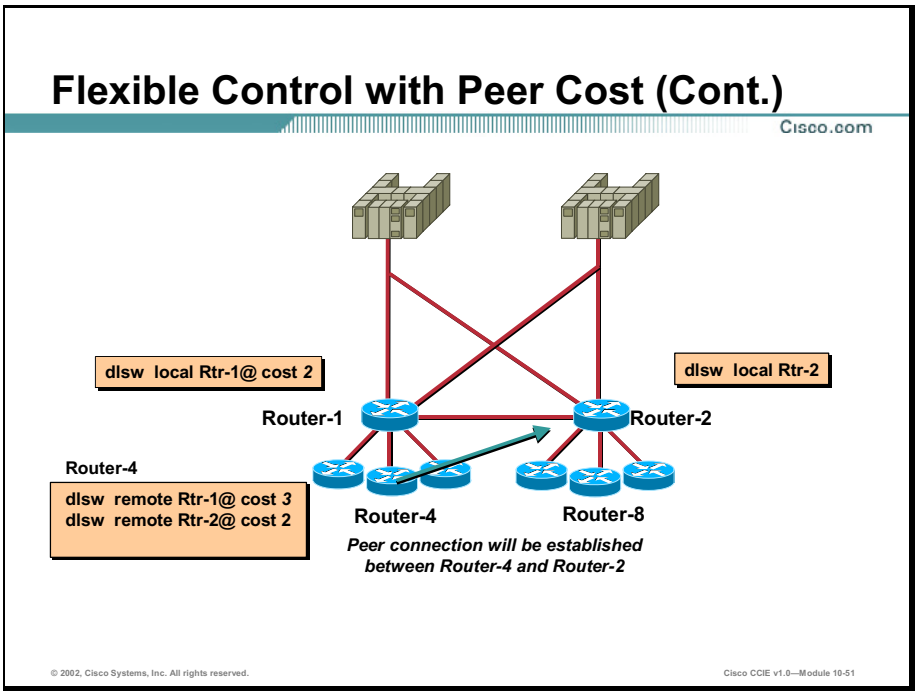

If costing is configured on the DLSw+ remote peer, then that cost overrides the default or configured cost value on the DLSw+ local peer.

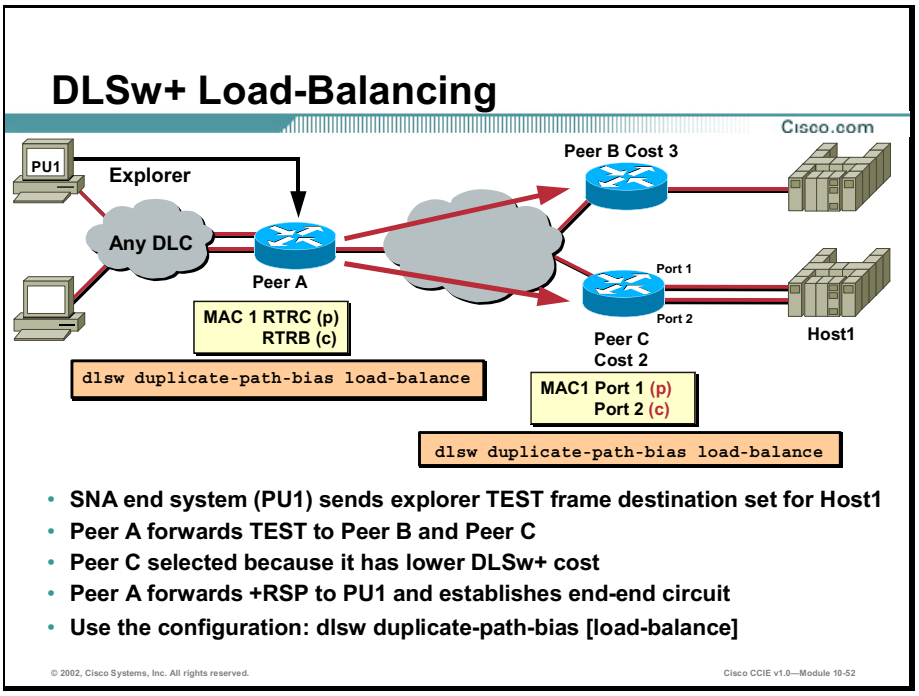

Use the configuration **dlsw duplicate-path-bias** [*load-balance*] for DLSw+ Load – Balancing.

- SNA end system Physical Unit (PU) 1 sends explorer TEST frame destination set for Host1
- Peer A forwards TEST to Peer B and Peer C
- Peer C is selected because it has lower DLSw+ Cost
- Peer A forwards +RSP to PU1 and establishes end-end circuit

DLSw+ Backup Peers

Cisco.com

Automatic recovery from loss of central site router

Connect to backup router only if primary fails

Automatic return to primary router when available:

- **New sessions—immediately**
- **Active sessions—configurable to prevent unnecessary session disruption**

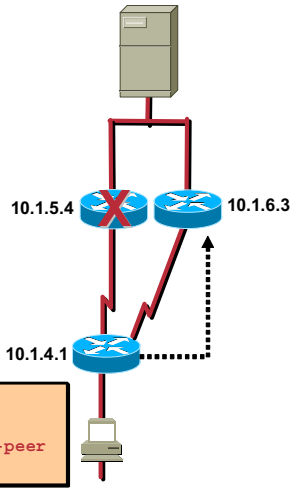


```
router(config)# dlsw local-peer peer-id 10.1.4.1
router(config)# dlsw remote-peer 0 tcp 10.1.5.4
router(config)# dlsw remote-peer 0 tcp 10.1.6.3 backup-peer
10.1.5.4 linger 0
```


Cisco CCIE v1.0—Module 10-53

DLSw+ backup capability allows connection to a backup DLSw+ peer (router) when the primary DLSw+ peer is lost. Loss of a DLSw+ peer disrupts all active sessions, but when the sessions are reactivated, they will automatically use the new DLSw+ peer, without operator intervention.

When the primary router recovers, all new sessions will be established using the primary peer. Existing sessions using the backup peer will continue unaffected until they terminate normally or until an optional linger time expires. The linger timer allows the administrator to move all sessions back to the primary peer as soon as the primary peer is active/stable. Optionally, the linger timer can be set longer to allow the network operator to warn the end users of an impending session disruption. The movement from the backup peer to the primary peer is disruptive to SNA sessions, but may be the preferred action if the backup peer resides at a disaster recovery site.

Summary

This section summarizes the key points discussed in this lesson.

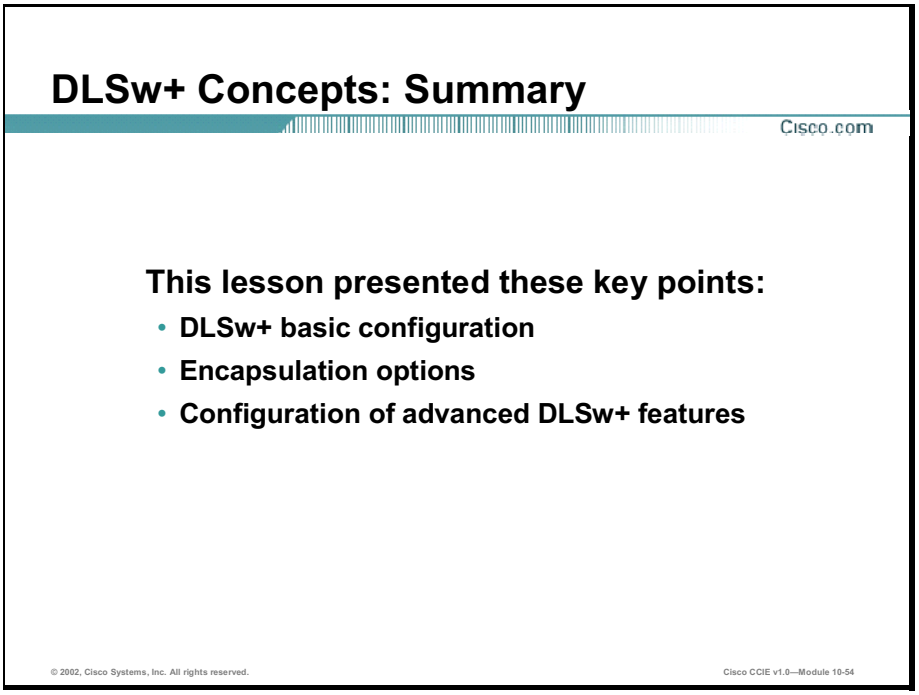

Next Steps

After completing this lesson, go to:

- DLSw+ Troubleshooting

References

For additional information, refer to these resources:

- Cisco Press *Cisco IOS Bridging and IBM Solutions*
- Designing DLSw+ Internetworks: http://www.cisco.com/univercd/cc/td/doc/cisintwk/idg4/nd2007.htm
- DLSw+ Peer Group Clusters: http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/120newft/120t/120t3/cluster.htm

Lesson Assessment (Quiz)

Q1) What are the minimum components of an Ethernet DLSw+ configuration? (Choose all that apply.)

A) Bridge-group

B) Pseudo bridge-group

C) Local Peer

D) Remote Peer

E) Virtual Ring

Q2) Which of the following is NOT a supported DLSw+ encapsulation type?

A) TCP

B) FST

C) Direct

D) STP

Q3) Which of the following permits a casual, any-to-any connection without the burden of configuring the connection in advance?

A) Border peers

B) Peer groups

C) On-demand peers

D) Explorer firewalls

Q4) What types of caching does DLSw+ use to reduce traffic? (Choose all that apply.)

A) Local

B) Remote

C) Border peer or group

D) On-demand

DLSw+ Troubleshooting

Overview

This lesson provides an overview of the tools used in troubleshooting problems that may be encountered in Data Link Switching Plus (DLSw+) networks.

Importance

The ability to quickly detect, diagnose, and resolve DLSw+ problems is critical to success on the CCIE lab.

Objectives

Upon completing this lesson, you will be able to:

- Use **show** commands to verify DLSw+ configurations
- Use **debug** commands to diagnose problems in a DLSw+ network

Learner Skills and Knowledge

To fully benefit from this lesson, you must have these prerequisite skills and knowledge:

- Passed the CCIE written qualification exam
- An understanding of bridging technologies, discussed in a previous lesson
- The ability to configure a DLSw+ network

Outline

This lesson includes these sections:

- Overview
- Show Commands
- Debug Commands
- Summary
- Lesson Assessment (Quiz)

Show Commands

This section examines various **show** commands that can be used to verify and diagnose problems in a Data Link Switching Plus (DLSw+) network.

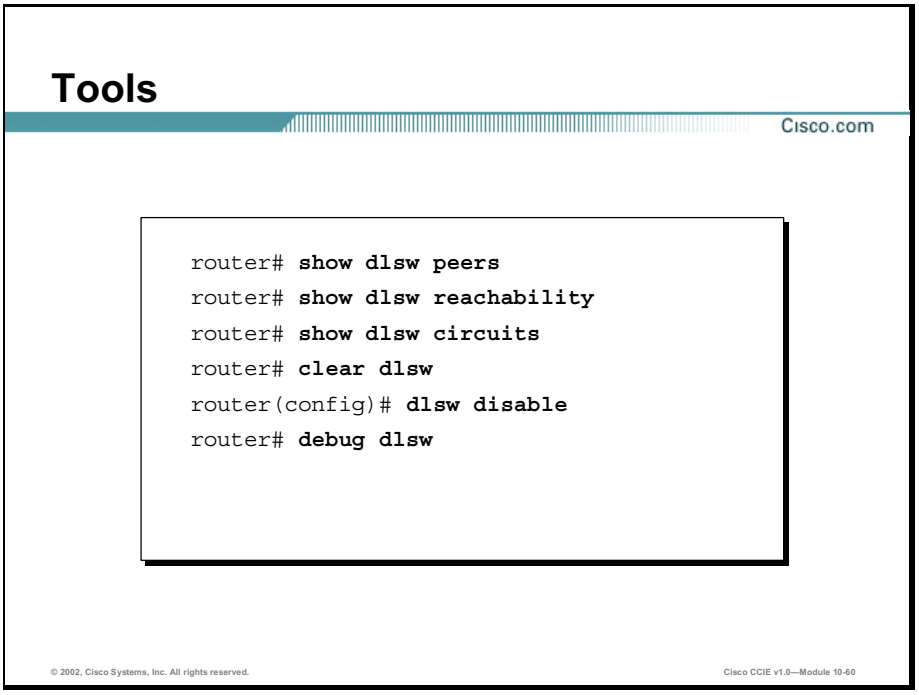

Cisco offers multiple **show** and **debug** commands for analyzing DLSw+ networks. Over the next several pages, selected **show** and **debug** commands will be examined.

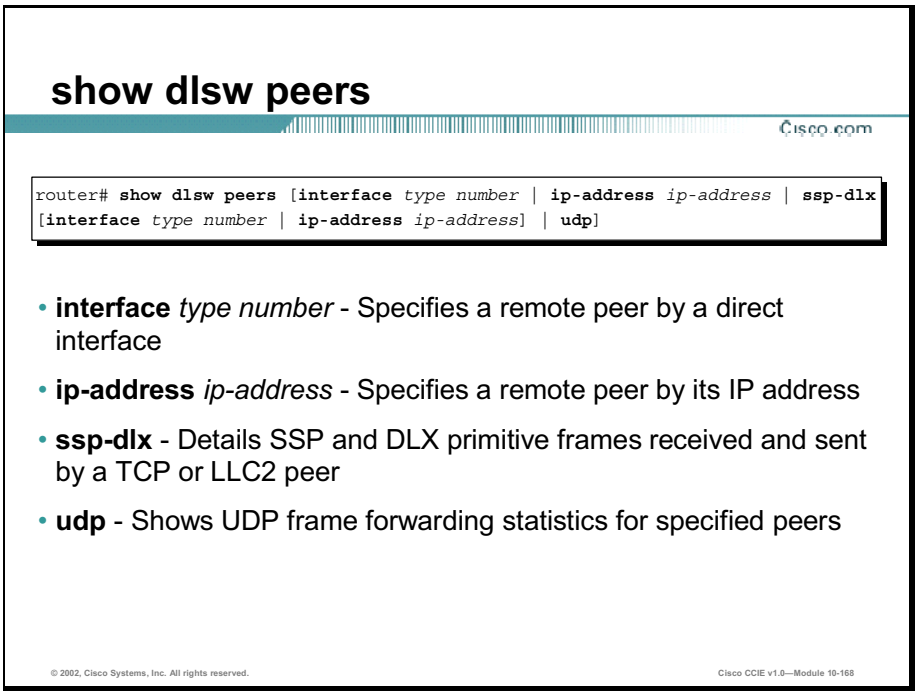

The **show dlsw peers** command allows you to show the status of remote peers. With the exception of local circuits, nothing happens in DLSw+ without remote peer connections. If the peer is not in CONNECT status, no data traffic can flow between end stations that are trying to traverse the peer connection.

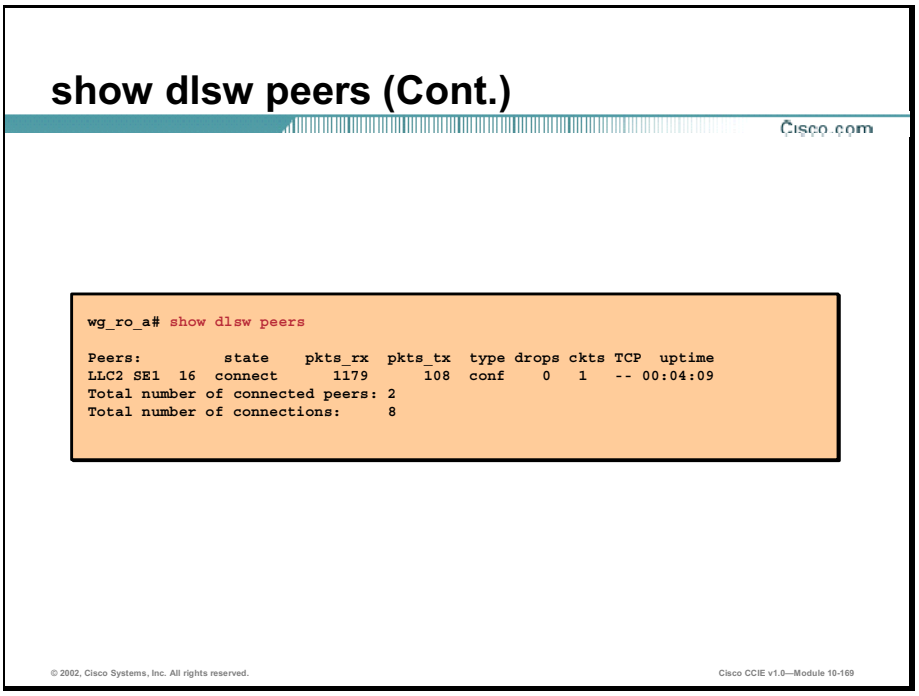

The following table details the various fields from the output of the **show dlsw peers** command.

Table 9-7: show dlsw peers Output Fields

Field	Description
`Peers`	Information related to the remote peer, including encapsulation type, IP address (if using FST or TCP), and interface number (if using direct encapsulation)
`tot-Q'd`	UDP packets that have been queued because of TCP congestion
`total-rx`	Number UDP packets received from the peer
`total-tx`	Number of UDP packets sent to the peer
`tot-retx`	Number of reachability resends (for example, DLSw+ retries, NQ_ex, and CUR_ex) when originally sent via UDP
`tot-drop`	Number of queued UDP packets that were dropped because of persistent TCP congestion
`curr-Q'd`	Number of current UDP packets queued because of TCP congestion
`TCP`	Number of packets currently on TCP output queue

Field	Description
state	State of the peer: ■ CONNECT – normal working peer ■ DISCONN – peer is not connected ■ CAP_EXG – capabilities exchange mode. Waiting for capabilities response ■ WAIT_RD – TCP write pipe (local port 2065) is open and peer is waiting for remote peer to open the read port (local port 2067). This field applies to only TCP peers ■ WAN_BUSY – TCP
pkts_rx	Number of received packets
pkts_tx	Number of sent packets
type	Type of remote peer: ■ conf – configured ■ prom – promiscuous ■ pod – peer on demand
drops	Number of drops completed by this peer. Reasons for the counter to increment: ■ WAN interface not up for a direct peer ■ DLSW tries to send a packet before the peer is fully connected (waiting for TCP event or capabilities event) ■ Outbound TCP queue full ■ FST sequence number count mismatch ■ Cannot get buffer to “slow switch” FST packet ■ CiscoBus controller failure on high end (cannot move packet from receive buffer to send buffer, or vice versa) ■ Destination IP address of FST packet does not match local peer-ID ■ WAN interface not up for an FST peer ■ No SRB route cache command configured ■ Madge ring buffer is full on low end systems (WAN feeding LAN too fast)
uptime	How long the connection has been established to this peer
ckts	Number of active circuits through this peer. This field applies only to TCP and LLC2 transport peer types
total number of connected peers	Total number of currently connected peers
total number of connections	Total number of active circuit connections

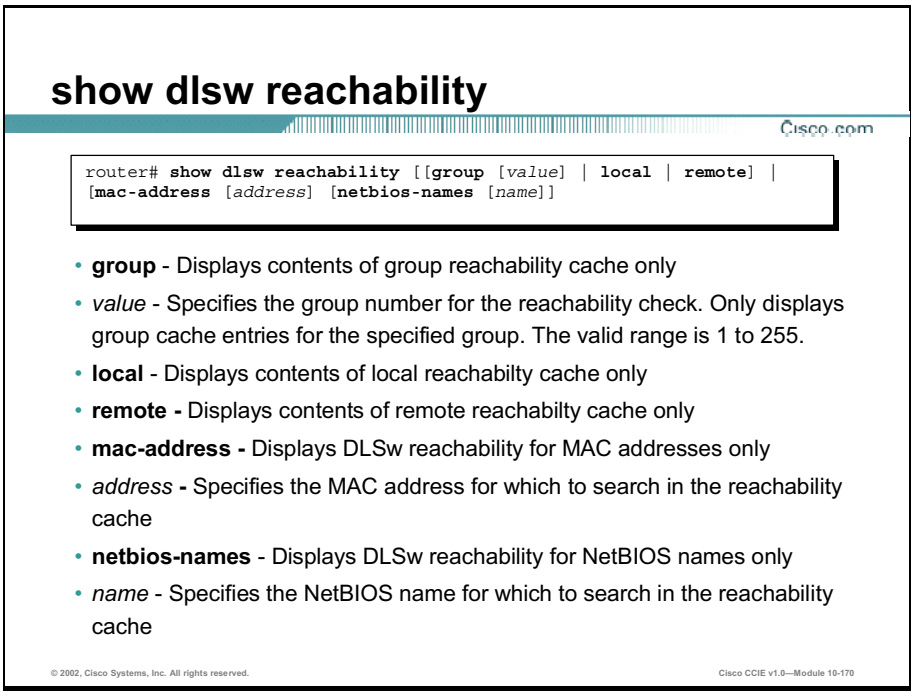

You can use the **show dlsw reachability** command to determine which System Network Architecture (SNA) or NetBIOS DLSw+ end stations a router has in its cache. DLSw+ checks the reachability cache when it is trying to initiate a session to determine if it recognizes the correct peer or port to use for this session. It also checks this cache when attempting to send traffic that is not session-based (that is, connectionless) across DLSw+. If DLSw+ does not know where a particular destination address is, it queries other peers that it knows. When it does learn how to reach a destination, DLSw+ keeps that information for a specific amount of time in an effort to reduce the broadcast traffic on the network.

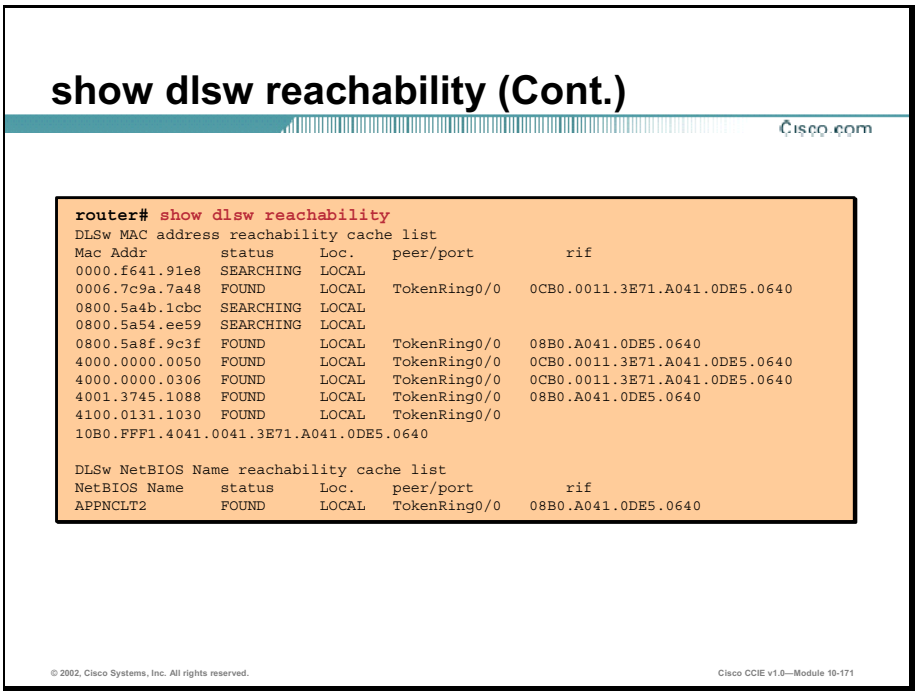

The following table defines the various fields from the output of the **show dlsw reachability** command.

Table 9-8: show dlsw reachability Output Fields

Field	Description
`Mac Addr`	MAC address of station being sought (destination MAC address of anureach_ex packet)
`NetBIOS Name`	NetBIOS name of station being sought (destination MAC address of NQ_ex packet)
`status`	Result of station search. The status can be one of the following: ■ FOUND – Station has recently sent a broadcast or responded to a broadcast. ■ SEARCHING – Router has sent broadcast to this station and is waiting for a response. ■ NOT_FOUND – Negative caching is on, and the station has not responded to queries. ■ UNCONFIRMED – Station is configured, but DLSw has not verified it. ■ VERIFY – Cache information is being verified because cache is going stale, or the user configuration is being verified.
`Loc.`	Location of station. LOCAL indicates that the station is on the local network. REMOTE indicates that the station is on the remote network.

Field	Description
peer/port	Peer/port number. If the Loc. Field lists a REMOTE station, the peer/port field indicates the peer through which the remote station is reachable. If the Loc. Field lists a LOCAL station, the peer/port field indicates the port through which the local station is reachable. For ports, the port number and slot number are given. Pxxx-Syyy denotes port xxx slot yyy. If the station is reachable through a bridge group, that is shown by Tbridge-xxx.
rif	Shows the RIF in the cache. This column applies only to LOCAL stations. If the station was reached through a medium that does not support RIFs (such as SDLC or Ethernet) then "—nofif—" is shown.

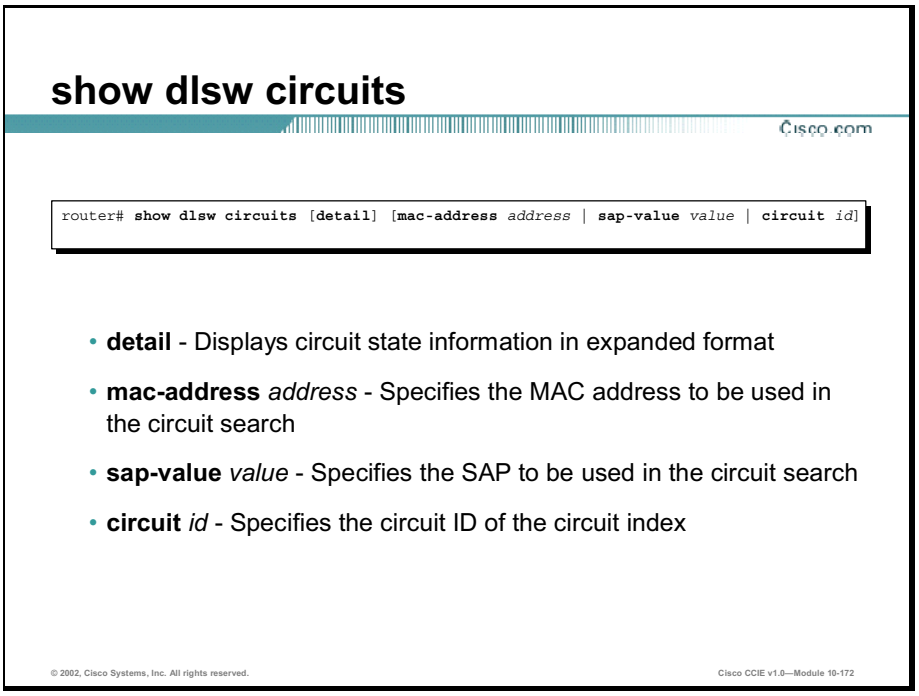

Use the **show dlsw circuits** command to display the state of all circuits involving a particular Media Access Control (MAC) address or Service Advertising Protocol (SAP) value.

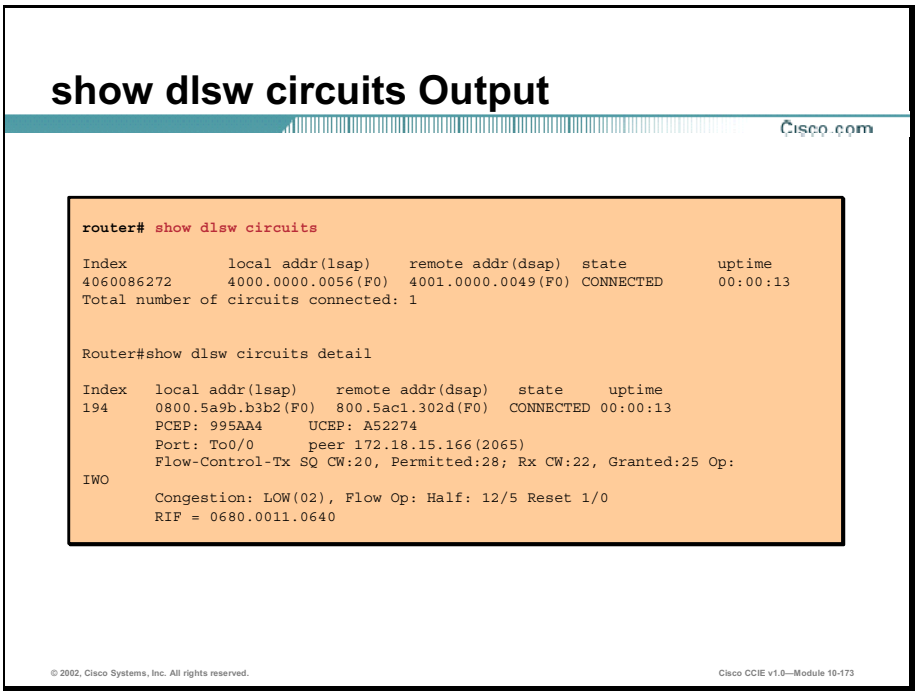

The following table defines the fields that appear in the output of the **show dlsw circuits** command.

Table 9-9: show dlsw circuits Output Fields

Field	Description
Index	Number the software uses to reference an individual circuit
local addr (lsap)	MAC address and SAP value used by end station closest to this DLSw+ peer
remote addr (dsap)	MAC address and SAP value used by end station that is across the peer connection (remote)
state	Indicates whether circuit has completed establishment
uptime	Length of time a circuit has been connected
total number of circuits connected	Number of total connected circuits. If a circuit has not completed connection, it will not show a value
PCEP, UCEP	Internal correlators used as labels for communication internal to the router between DLSw+ and LLC2, SDLC, or QLLC
Port	Local port over which this circuit has been established or DLSw interface to the bridge group
Flow Control (Tx and Rx)	Reports DLSw+ flow control windows as described in Section 8 of RFC 1795
SQ	Two flags indicating congestion toward the remote peer
S	Data flow from the local station has been stopped. This results in LLC2 or SDLC sending RNR frames
Q	Data frames are being queued for transport to the remote peer

Field	Description
CW	Current pacing window. See RFC 1795
Permitted	Packet counter for tx. See RFC 1795
Granted	Packet counter for rx. See RFC 1795
Op	Next flow indicator (FCI) that will be sent to the remote peer. See RFC 1795
Congestion	Data flow indicator from router to station is congested. Values are low, medium, high, and max
Flow Op	Amount of Reset Window Operator and Half Window Operator being sent or received. See RFC 1795
RIF	Routing Information Field used over the local port for data traversing this circuit (if appropriate)

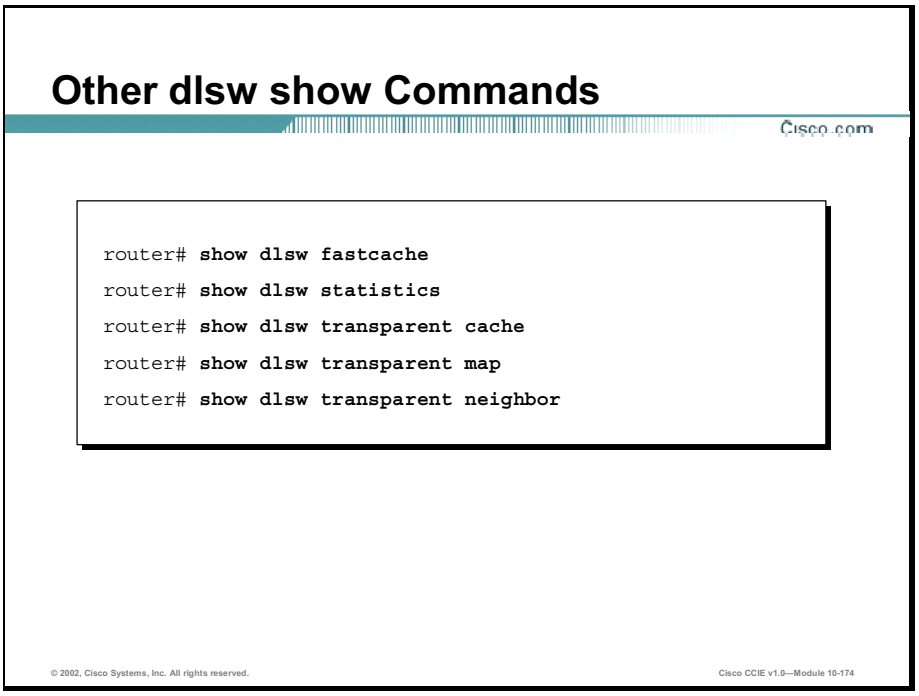

Below is a collection of additional commands for DLSw+ verification and troubleshooting:

- **show dlsw fastcache** - Displays the fast cache for Fast-Sequenced Transport (FST) and direct-encapsulated peers
- **show dlsw statistics** - Displays the number of frames that have been processed in the local, remote, and group cache
- **show dlsw transparent cache** - Displays the master circuit cache for each transparent bridged domain
- **show dlsw transparent map** - Displays MAC address mappings on the local router and any mappings for which the local router is acting as backup for a neighbor peer
- **show dlsw transparent neighbor** - Displays DLSw+ neighbors in a transparent bridged domain

Debug Commands

For troubleshooting purposes, Cisco provides multiple options for the **debug dlsw** command.

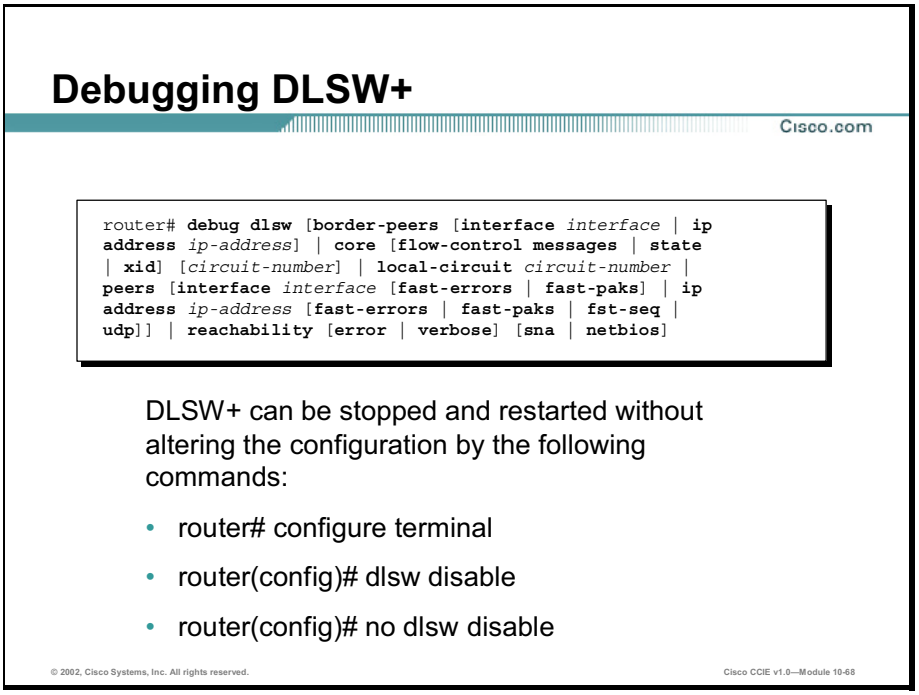

The following table defines the syntax parameters of the **debug dlsw** command.

Table 9-10: debug dlsw Syntax Parameters

Syntax	Description
`border-peers`	Enables debugging output for border peer events
`interface interface`	Specifies a remote peer to debug by a direct interface
`core`	Enables debugging output for DLSw core events
`flow-control`	Enables debugging output for congestion in the WAN or at the remote end station
`messages`	Enables debugging output of core messages – specific packets received by DLSw either from one of its peers or from a local medium via the Cisco link services interface
`state`	Enables debugging output for state changes on the circuit
`xid`	Enables debugging output for the exchange identification state machine
`circuit-number`	Specifies the circuit where you want core debugging output to reduce the amount of debugging output
`local-circuit circuit-number`	Enables debugging output for circuits performing local conversion. Local conversion occurs when both the input and ouput data-link connections are on the same local peer and no remote peer exists

Syntax	Description
peers	Enables debugging output for peer events
fast-errors	Debugs errors for fast-switched packets
fast-paks	Debugs fast-switched packets
fst-seq	Debugs FST sequence numbers on fast switched packets
udp	Debugs UDP packets
reachability	Enables debugging output for reachability events (explorer traffic). If no options are specified, event-level information is displayed for all protocols
error | verbose	Specifies how much reachability information you want displayed. The verbose keyword displays everything, including errors and events. The error keyword displays error information only. If no option is specified, event-level information is displayed
sna | netbios	Specifies that reachability information be displayed for only SNA or NetBIOS protocols. If no option is specified, information for all protocols is displayed

Summary

This section summarizes the key points discussed in this lesson.

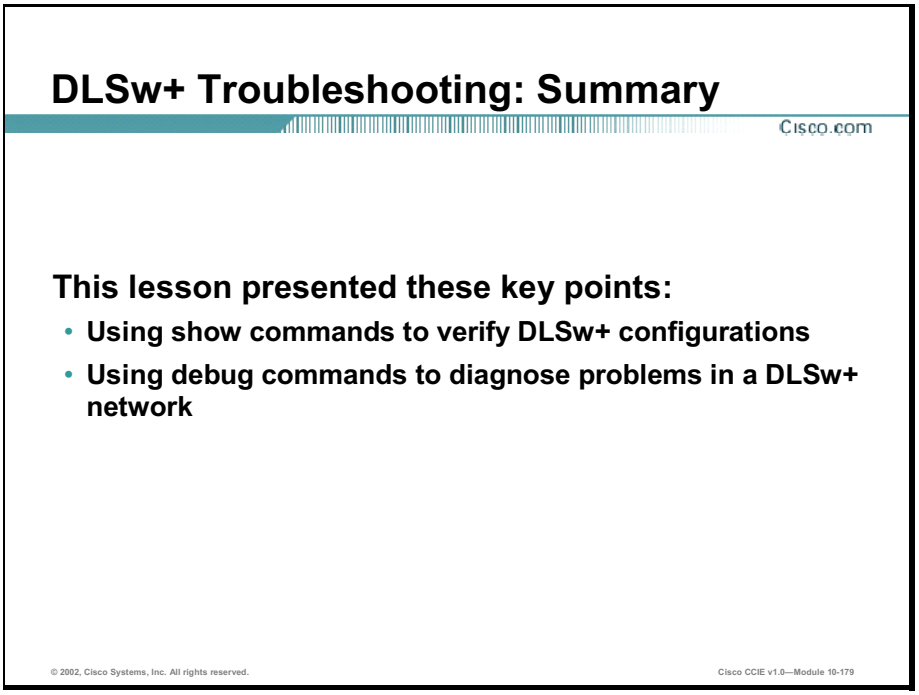

Next Steps

After completing this lesson, go to:

- Multicasting and IP Services

References

For additional information, refer to these resources:

- Troubleshooting DLSw: http://www.cisco.com/warp/public/697/dlswts1.html
- Using SHOW and DEBUG Commands: http://www.cisco.com/warp/public/cc/pd/ibsw/ibdlsw/tech/dls9_rg.htm
- Troubleshooting DLSw Circuit Connectivity: http://www.cisco.com/warp/public/697/dlswts4.html
- Cisco Press *Cisco IOS Bridging and IBM Network Solutions*

Lesson Assessment (Quiz)

Q1) Which of the following commands allows you to show the status of remote peers?

A) **show dlsw status**

B) **show dlsw peers**

C) **show dlsw peer status**

D) **show status**

Q2) Which of the following commands will let you determine which SNA or NetBIOS DLSw+ end stations a router has in its cache?

A) **show dlsw peers**

B) **show dlsw cache**

C) **show dlsw reachability**

D) **show dlsw circuits**

Q3) Which of the following commands displays the state of all circuits involving a particular MAC address or SAP value?

A) **show dlsw circuits**

B) **show dlsw reachability**

C) **show dlsw peers**

D) **show dlsw cache**

Q4) What is the normal state of a DLSw+ circuit when it is successfully connected?

A) MULTIPLE-FRAMES-ESTABLISHED

B) UP

C) CONNECTED

D) PEER-PEER

Q5) Which of the following commands enables UDP debugging for a remote peer, with an IP address of 1.1.1.6?

A) debug dlsw remote-peer 1.1.1.6 udp

B) debug dlsw peer ip-address 1.1.1.6 udp

C) debug dlsw remote udp 1.1.1.6

D) debug dlsw udp remote ip-address 1.1.1.6

10

Multicasting and IP Services

Overview

This module covers configuration of Internet Protocol (IP) Multicast routing protocols and their interoperation with Cisco Switches.

Upon completing this module, you will be able to:

- Describe multicast operation using Protocol Independent Multicast-Dense Mode (PIM-DM), Protocol Independent Multicast-Sparse Mode (PIM-SM) and Internet Group Multicast Protocol (IGMP)
- Enable multicast routing
- Configure Network Time Protocol (NTP)
- Configure and verify Network Address Translation (NAT)
- Configure Hot Standby Router Protocol (HSRP)
- Configure Dynamic Host Configuration Protocol (DHCP)

Outline

The module contains these lessons:

- Multicasting Configuration
- Network Time Protocol
- Network Address Translation
- Hot Standby Routing Protocol

- Dynamic Host Configuration Protocol

Multicasting Configuration

Overview

Multicast routing is a rather complicated topic, but its configuration is straightforward. This lesson will discuss multicast routing commands for Protocol Independent Multicast – Dense Mode (PIM-DM) and Protocol Independent Multicast – Sparse Mode (PIM-SM), Cisco Group Message Protocol (CGMP), and Internet Group Management Protocol (IGMP) snooping.

Importance

Multicast technologies are a critical aspect of the CCIE lab. Configuration of Internet Protocol (IP) Multicast routing protocols and their interoperation with Cisco Switches is essential for one to be successful on the exam.

Objectives

Upon completing this lesson, you will be able to:

- Enable multicast routing on a Cisco router
- Define what PIM protocol will be utilized
- Configure Cisco switches to be Multicast flow aware

Learner Skills and Knowledge

To fully benefit from this lesson, you must have these prerequisite skills and knowledge:

- Passed the CCIE written qualification exam
- Completed the Building Cisco Multilayer Switched Networks (BCMSN) course or have the equivalent knowledge

Outline

This lesson includes these sections:

- Overview
- Router Configuration
- Multicast Routing Protocols
- Summary
- Lesson Assessment (Quiz)

Router Configuration

Cisco Routers by default in Internetwork Operating System version 12.0 (IOSv12.0) listen to Internet Group Management Protocol version 2 (IGMPv2) on all Internet Protocol (IP) addressed interfaces. The router is also aware of version 1 devices and will apply appropriate timers to maintain proper multicast functions.

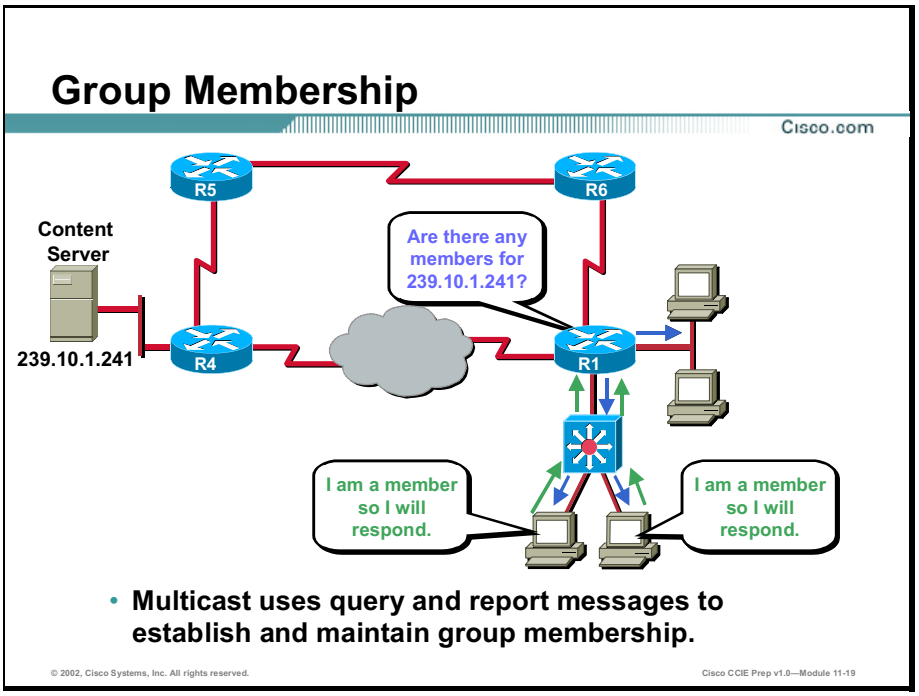

According to the IGMPv1 specification, one multicast router per Local Area Network (LAN) must periodically transmit Host Membership Query messages to determine which host groups have members on the querier's directly attached networks. IGMP query messages are addressed to the all-host group (224.0.0.1) and have an IP Time To Live (TTL) equal to one. This TTL ensures the query messages sourced from a router are transmitted onto the directly attached network but are not forwarded by any other multicast routers.

When the end station receives an IGMP query message, the end station responds with a host membership report for each group into which the end station belongs.

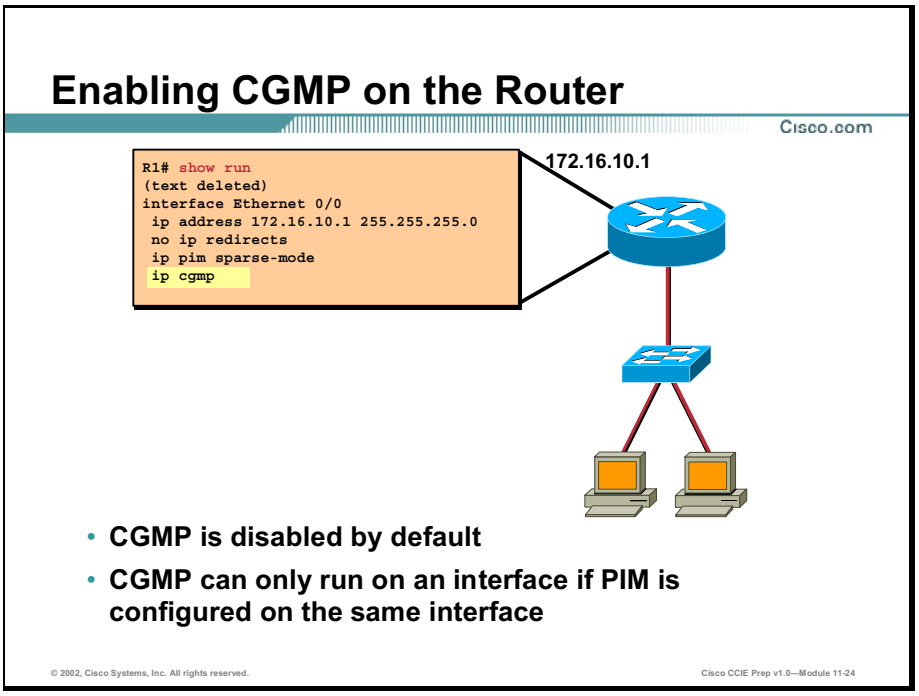

CGMP can only run on an interface if Protocol Independent Multicast (PIM) is configured on the same interface. CGMP is disabled by default. To enable CGMP on the router, enter the following command in the interface configuration mode.

```
R1(config-if)# ip cgmp
```

This command enables the CGMP for IP multicast on a router. This command triggers a CGMP Join Message.

The running configuration indicates if a specific router interface has been configured for CGMP.

```
R1# show run
Building configuration...
Current configuration:
!
(text deleted)
interface Ethernet 0/0
 ip address 172.16.10.1 255.255.255.0
 no ip redirects
 ip pim sparse-mode
 ip cgmp
```

When a **no ip cgmp** command is issued, a triggered CGMP Leave Message for group 0000.0000.0000 is sent with the router's Media Access Control (MAC) address on the interface.

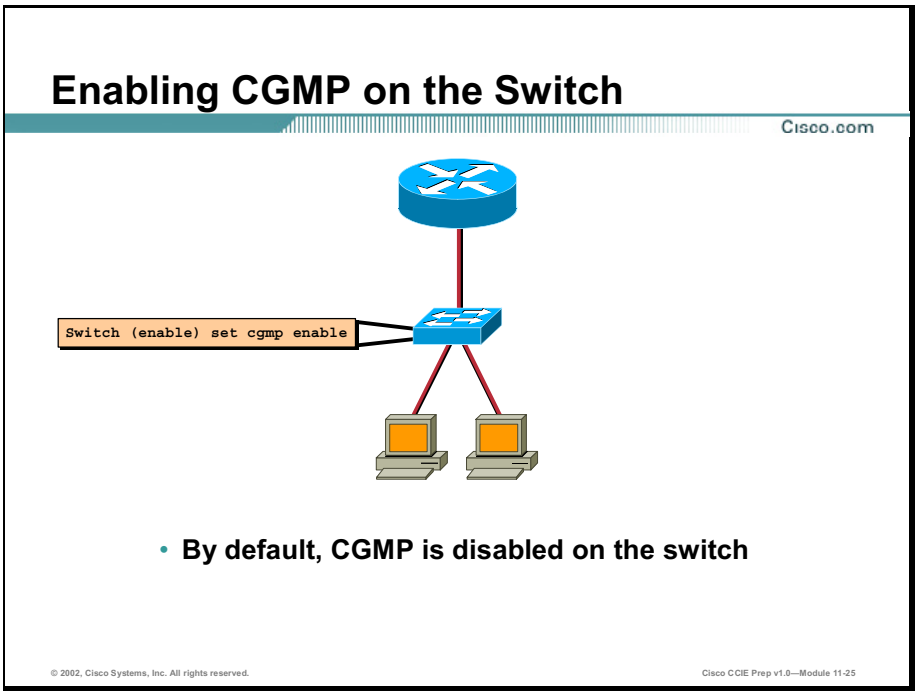

To enable CGMP on a switch, enter the following command in EXEC mode.

```
Switch (enable) set cgmp enable
```

Note	IGMP snooping must be disabled before you can enable CGMP.

The running configuration indicates if the specific switch has been configured for CGMP.

```
Switch (enable) show config
(text deleted)
!
#cgmp
set cgmp enable
set cgmp leave disable
```

To disable CGMP on a switch, enter the **set cgmp disable** command.

Configuring CGMP on the switch allows IP multicast packets to be switched only to those ports that have IP multicast clients. Directing multicast traffic only to those user segments that have interested clients reduces the consumption of network bandwidth by not propagating IP multicast traffic throughout the broadcast domain. CGMP on a switch also reduces management and resource overhead by not requiring a separate Virtual LAN (VLAN) for each multicast group in the switched network to separate traffic.

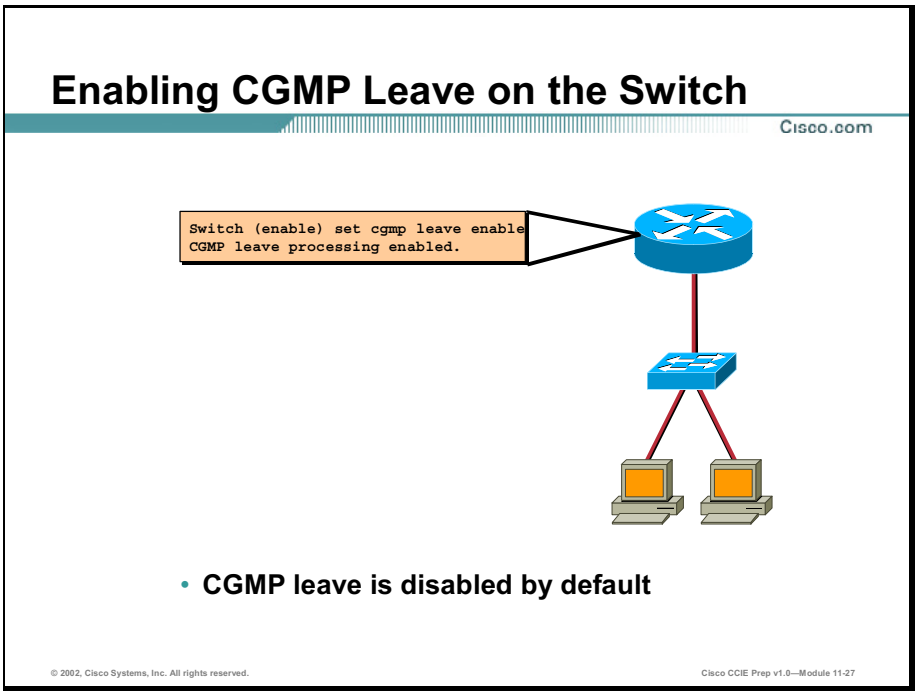

To enable the CGMP leave function on the switch, enter the following command in privileged EXEC mode.

```
Switch (enable) set cgmp leave
```

The **show cgmp leave** command provides verification that CGMP fast-leave has been configured. This command can be entered in either user or privileged EXEC mode.

```
Switch (enable) show cgmp leave
CGMP: enabled
CGMP leave: enabled
```

To disable the CGMP fast-leave function, enter the **set cgmp leave disable** command.

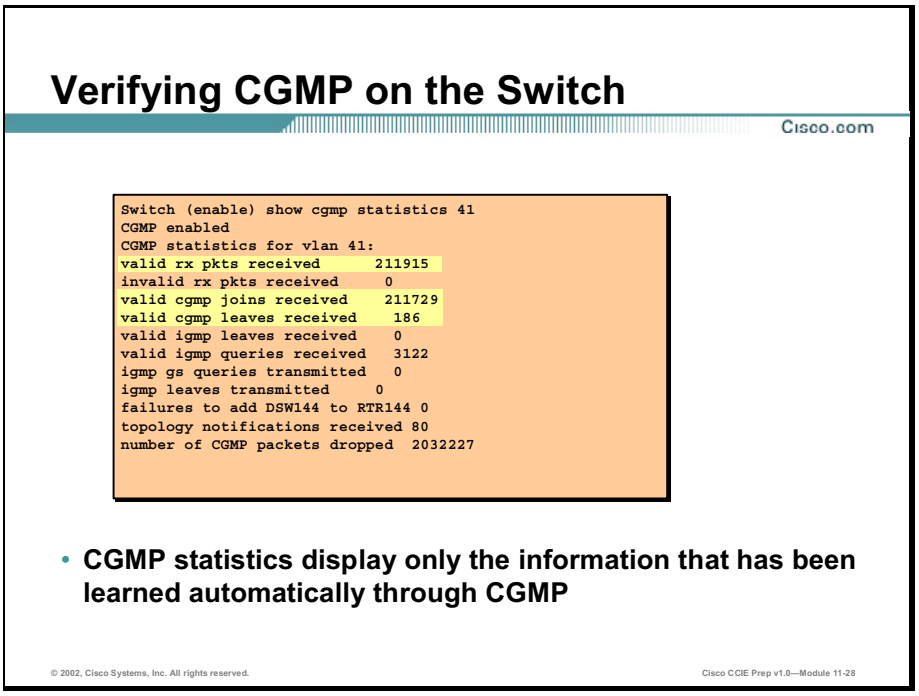

Several commands are available to help you verify the CGMP configuration on the switch. These commands allow you to view the relationships between VLANs, CGMP-enabled routers, and switch ports.

Note	Refer to the “Configuring Multicast Services” section in the *Catalyst 5000 Series Software Configuration Guide (4.3)* for a complete list of these commands.

The following two commands are useful in helping you verify the CGMP configuration on the switch.

The **show cgmp statistics** *vlan* command displays the ongoing CGMP activity for a designated VLAN.

The **show multicast group cgmp** *vlan command* displays only the multicast router information that has been learned automatically through CGMP. The results of this command show the multicast group MAC addresses associated with a specific VLAN and the ports associated with those groups.

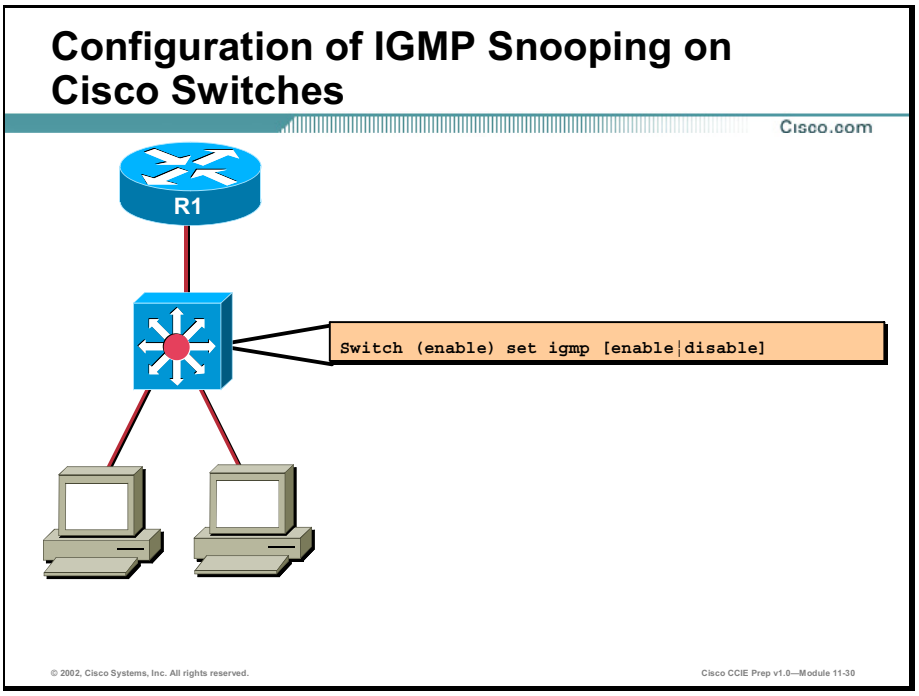

IGMP snooping, as implied by the name, is a feature that allows the switch to "listen in" on the IGMP conversation between hosts and routers. When a switch hears an IGMP report from a host for a given multicast group, the switch adds the host's port number to the Global Destination Address (GDA) list for that group. And, when the switch hears an IGMP leave, it removes the host's port from the Content-Addressable Memory (CAM) table entry.

`Switch (enable) set igmp` `[enable|disable]`

Use this command to enable IGMP Snooping on a Catalyst switch.

Multicast Routing Protocols

This section covers multicast routing protocols.

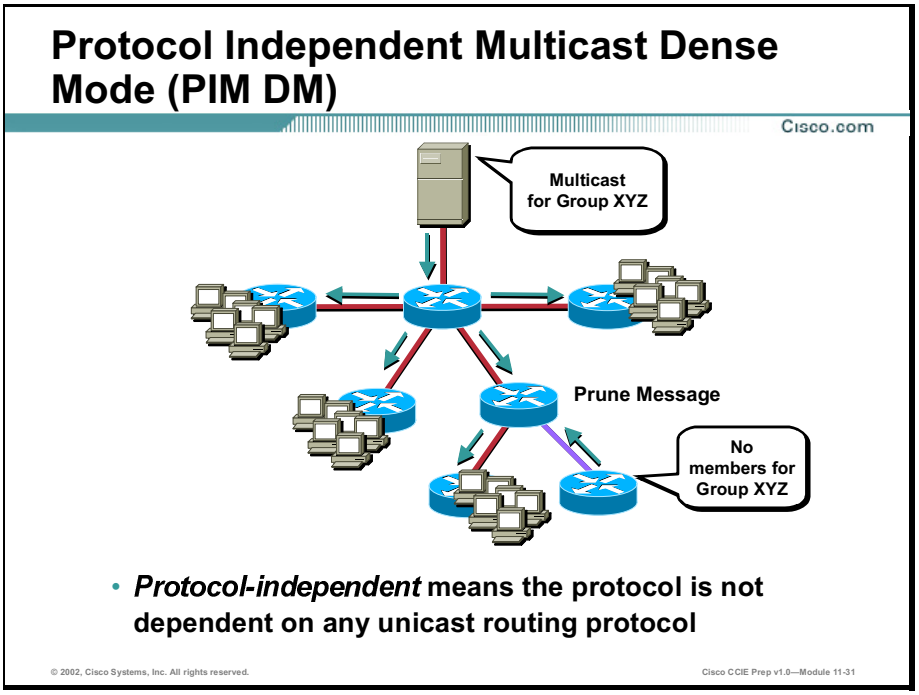

Protocol Independent Multicast (PIM) operates independent of the routing protocol, but it uses the routing protocol as the foundation for building the tree. If the routing protocol has not converged or does not know of the network, nor will PIM.

Protocol Independent Multicast Dense Mode (PIM-DM) works best when there are numerous members belonging to each multimedia group. PIM floods the multimedia packet out to all routers in the network and then prunes routers that do not support members of that particular multicast group.

PIM-DM is most useful when:

- Senders and receivers are in close proximity to one another
- There are few senders and many receivers
- The volume of multicast traffic is high
- The stream of multicast traffic is constant

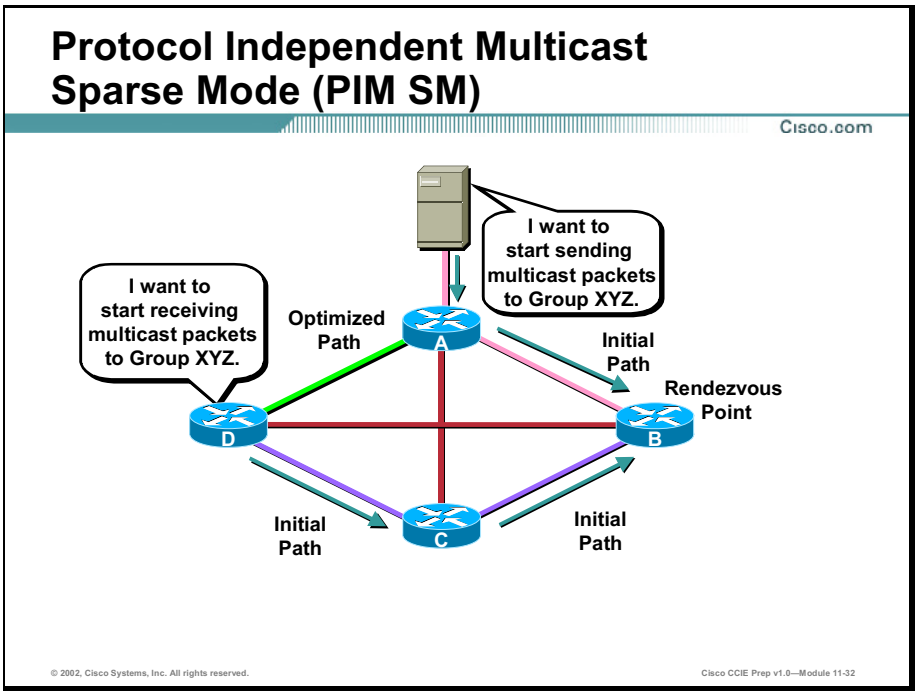

Protocol Independent Sparse Mode (PIM-SM) routing is based on the assumption that the multicast group members are sparsely distributed throughout the network and bandwidth is not necessarily widely available.

It is important to note that sparse mode does not imply that the group has few members, just that they are widely dispersed. In this case, flooding would unnecessarily waste network bandwidth and could cause serious performance problems. Therefore, sparse mode multicast routing protocols must rely on more selective techniques to set up and maintain multicast trees. Sparse-mode protocols begin with an empty distribution tree and add branches only as the result of explicit requests to join the distribution.

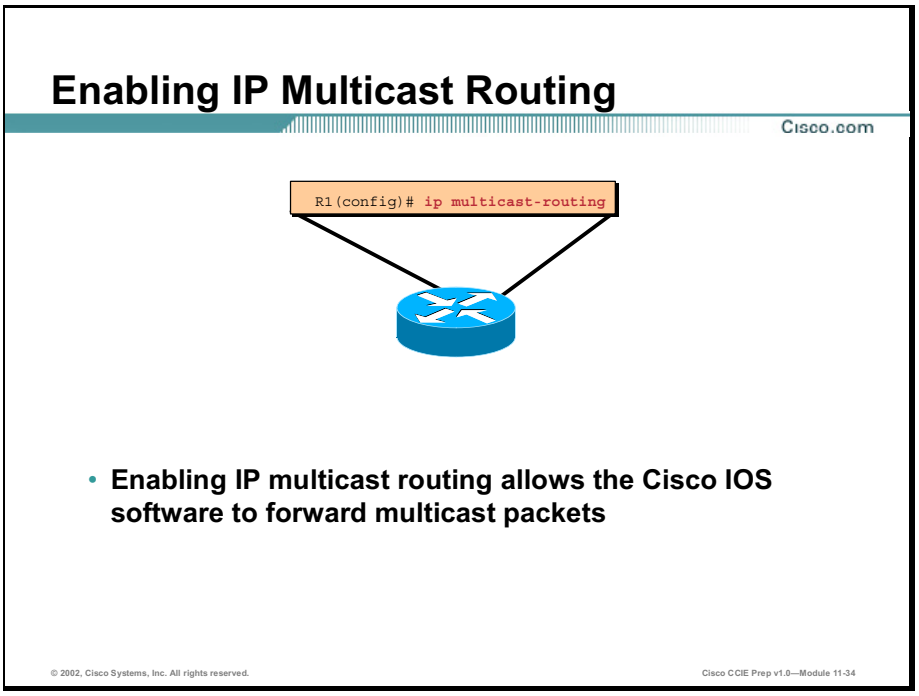

Enabling IP multicast routing allows the Cisco IOS software to forward multicast packets. To enable IP multicast routing on the router, enter the following command in global configuration mode.

```
R1(config)# ip multicast-routing
```

To disable IP multicast routing, enter the **no ip multicast-routing** command.

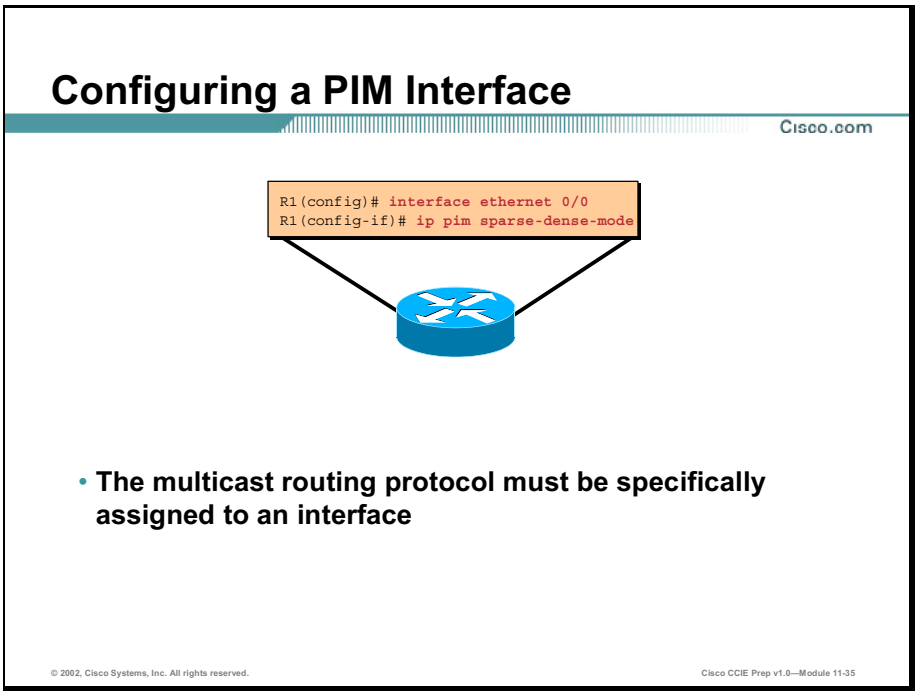

Routers forward multicast data on a per-interface basis. Therefore, the multicast routing protocol must be specifically assigned to an interface. To enable PIM on an interface, enter the following command in interface configuration mode.

```
R1(config-if)# ip pim {dense-mode | sparse-mode | sparse-dense-mode}
```

- **dense-mode:** Enables dense mode of operation. Dense mode is used when all routers in the network will need to distribute multicast traffic for each multicast group.
- **sparse-mode:** Enables sparse mode of operation. Sparse mode is used when relatively few routers in the network will be involved in each multicast.
- **sparse-dense-mode:** Treats the interface in the mode in which the group operates.

In sparse-dense mode, the interface is treated as dense mode if no rendezvous point is detected; the interface is treated as sparse mode if a rendezvous point is detected.

To disable PIM on the interface, enter the **no ip pim** command.

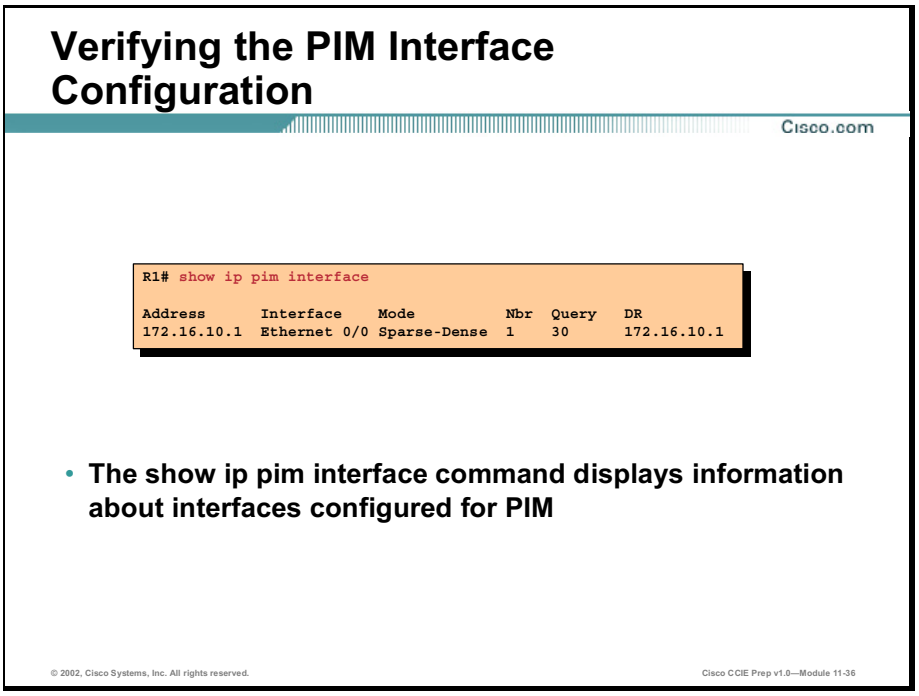

To display information about interfaces configured for PIM, enter the following command in EXEC mode.

R1# show ip pim interface [*type number*] [count]

Table 10-1

Command	Description
type	(Optional) Indicates the type of interface
number	(Optional) Indicates the number of the interface
count	(Optional) Displays the number of packets received and sent out the interface

Running this command displays the following items:

- The IP address of the next-hop router
- The interface type (On an Route Switch Module (RSM), this value will be the virtual LAN (VLAN) designation.)
- The PIM mode configured on each interface
- The number of PIM neighbors that have been discovered through this interface
- The frequency, in seconds, of PIM router query messages
- The IP address of the designated router on the LAN

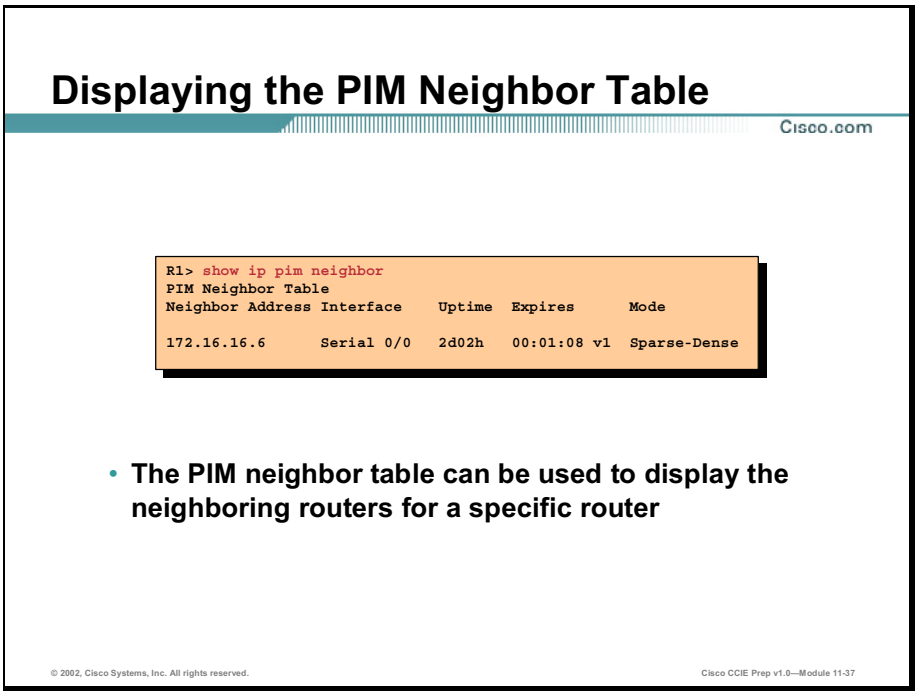

The Cisco IOS software can be used to discover the PIM neighbors in the network. To display a table of the neighboring routers from a specific router, enter the following command in privileged EXEC command mode.

R1>show ip pim neighbor *type number*

Table 10-2

Command	Description
type	(Optional) Indicates the type of interface
number	(Optional) Indicates the number of the interface

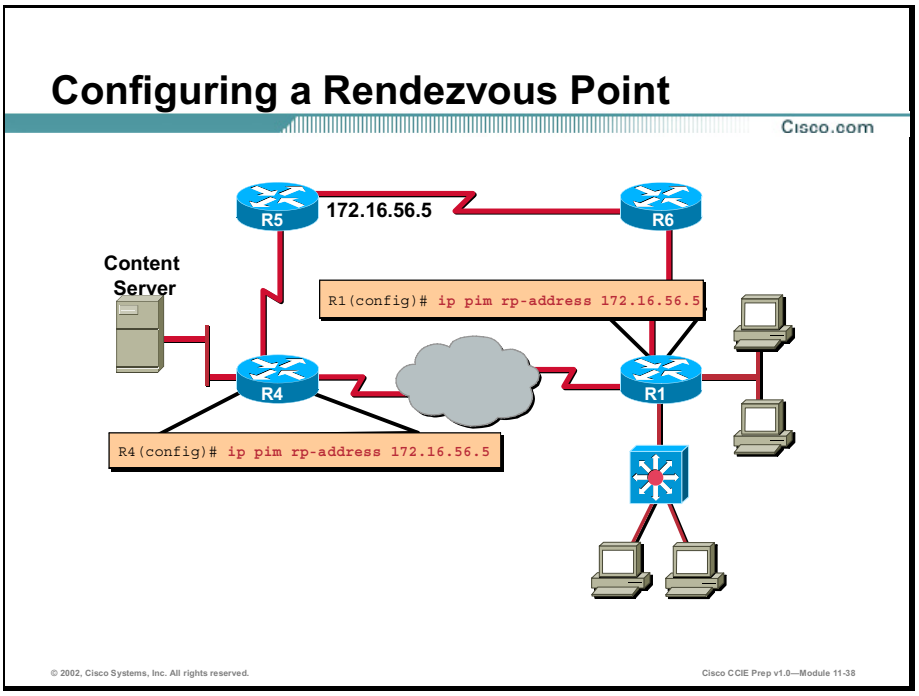

If you configure PIM to operate in sparse mode, you must also choose one or more routers to act as rendezvous points.

The IP address of the rendezvous points must be configured on leaf routers. Leaf routers are those routers that are directly connected either to a multicast group member or to a sender of multicast messages.

The Rendezvous Point (RP) address is used in either of the following scenarios:

- First-hop routers use the RP address to send PIM register messages on behalf of a host sending a packet to the group.
- Last-hop routers use the RP address to send PIM join or prune messages to the RP to inform it about group membership.

The RP does not need to know it is an RP. A PIM router can be an RP for more than one group. A group can have more than one RP. The conditions specified by the access list determine for which groups the router is an RP.

To designate the RP address on a leaf router, enter the following command in global configuration mode.

```
R4(config)# ip pim rp-address ip-address [group-access-list-number] [override]
```

Table 10-3

Command	Description
ip-address	Designates the IP address of a router to be a PIM RP. This is a unicast IP address in four-part, dotted notation.

group-access-list-number	(Optional) Number of an access list that defines for which multicast groups the RP should be used. This is a standard IP access list. The number can be from 1 to 100.
override	(Optional) Indicates that if there is a conflict between the RP configured with this command and one learned by auto-RP, the RP configured with this command prevails.

To remove an RP address, enter the **no ip pim rp-address** *ip-address* [*group-access-list-number*] command.

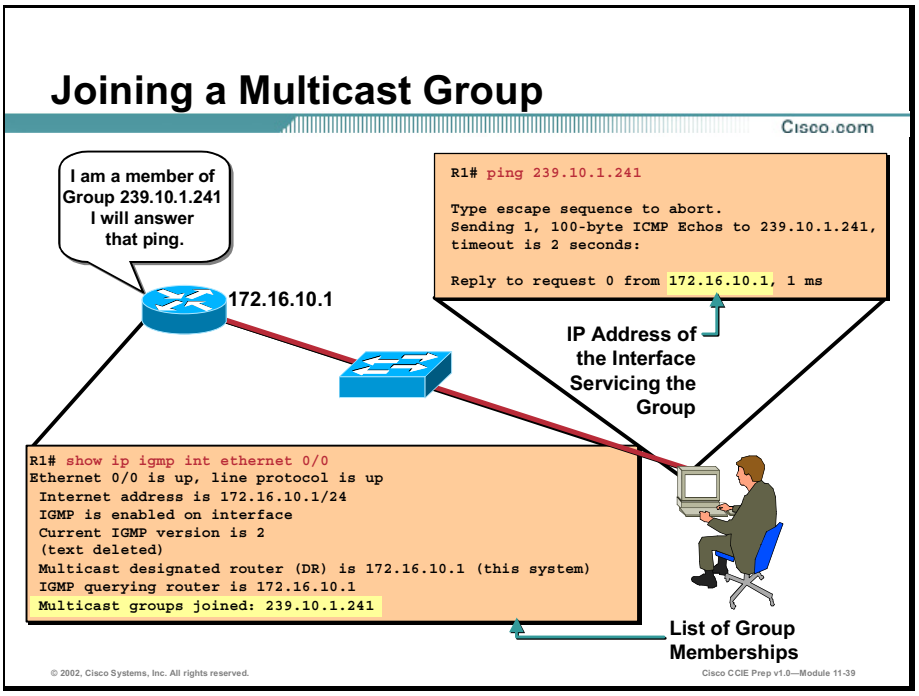

Being a member of a multicast group is useful in determining multicast reachability in a network. If a router is configured to be a member of a specific multicast group, that router can respond to commands, such as ping and Internet Control Message Protocol (ICMP) echo requests, addressed to that group. A group member router can also participate in multicast Cisco IOS trace route actions.

To add the router to a multicast group, enter the following command in interface configuration mode.

```
R1(config-if)# ip igmp join-group group-address
```

The *group-address* is the address of the multicast group.

Issuing a **ping** command specifying that multicast group address will cause all routers in that group to respond.

To cancel a router's membership in a multicast group, enter the **no ip igmp join-group** *group-address* command.

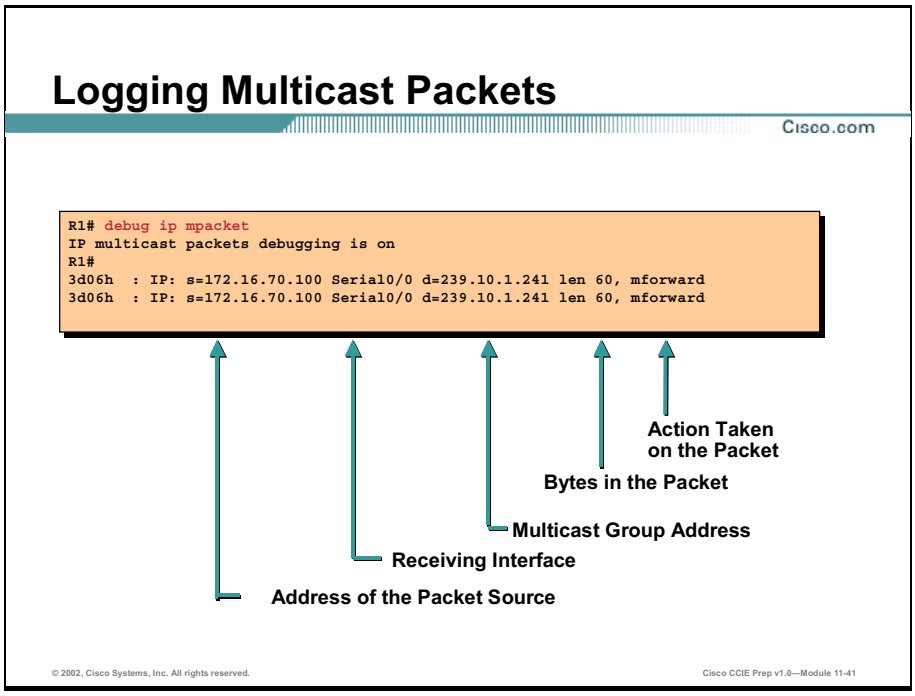

To log all IP multicast packets received and transmitted by a router, enter the following command in EXEC mode.

Caution	This command generates a substantial amount of messages and can impact router performance.

R1# debug ip mpacket [detail] [*acl*] [*group*]

Table 10-4

Variable	Definition
detail	(Optional) Monitors IP header information as well as MAC address information
acl	(Optional) Monitors only those multicast packets from sources described by the access list
group	(Optional) Monitors multicast packets generated by a single group

Summary

This section summarizes the key points discussed in this lesson.

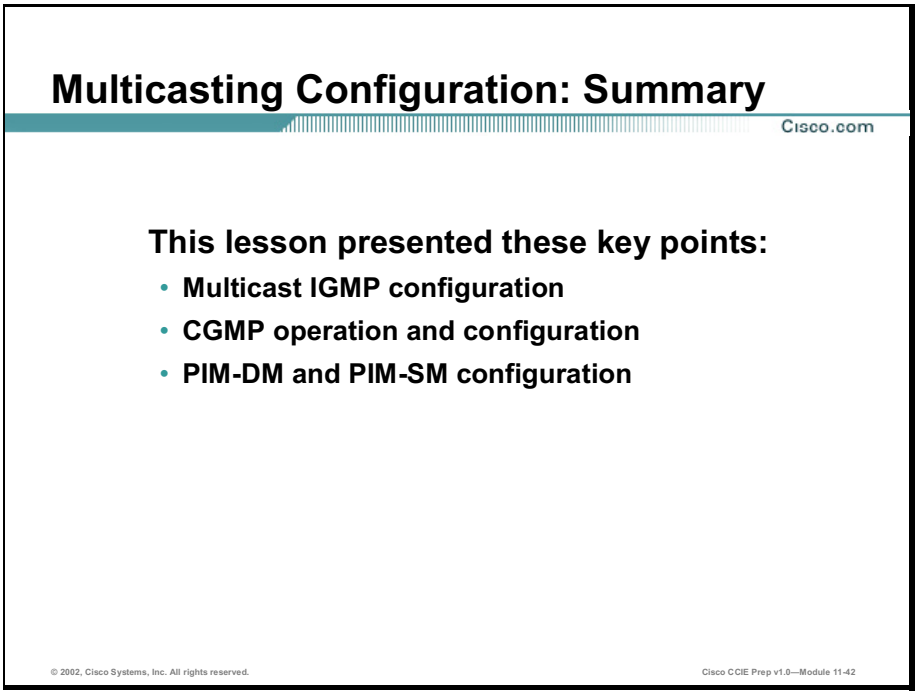

Next Steps

After completing this lesson, go to:

- Network Time Protocol

References

For additional information, refer to these resources:

- www.cisco.com
- *Routing TCP/IP, Volume II* by Jeff Doyle and Jennifer Dehaven Carroll
- *Developing IP Multicast Networks, Volume I* by Beau Williamson
- RFC 1112, 2236

Lesson Assessment (Quiz)

Q1) After configuring ip multicast routing, which steps must be taken to complete configuration of PIM-DM?

A) router(config-if)# **ip pim dense-mode** (on all interfaces)

B) router(config-if)# **ip pim dense-mode** (only interfaces used in multicasting)

C) router(config-if)# **ip pim dense-mode** and

router(config-if)# **ip igmp** (on all interfaces)

D) No additional configuration

Q2) CGMP requires _____.

A) a router

B) a Cisco Router

C) a CGMP enabled Cisco Router

D) nothing to run

Q3) True or False: IGMP snooping and CGMP are mutually exclusive.

A) True

B) False

Q4) Which command will inform you of the source of multicast packets?

A) **show ip mroute**

B) **show ip mroute source**

C) **show ip multicast source**

D) **debug ip packet**

Q5) What does (*,G) in the routing table tell about the multicast protocol that is running?

A) DVMRP

B) PIM-DM

C) IGMP

D) PIM-SM

Network Time Protocol

Overview

Network Time Protocol (NTP) allows for consistent time to be applied across a network. A central time server distributes time synchronization across the network.

Importance

NTP is utilized in the CCIE exam to aid in management of the devices.

Objectives

Upon completing this lesson, you will be able to:

- Describe NTP concepts
- Utilize NTP configuration

Learner Skills and Knowledge

To fully benefit from this lesson, you must have these prerequisite skills and knowledge:

- Passed the CCIE written qualification exam
- Completed the Interconnecting Cisco Network Devices (ICND) course or have the equivalent knowledge

Outline

This lesson includes these sections:

- Overview
- Network Time Protocol Concepts
- Basic Configuration
- Authentication Configuration
- Timezones
- Summary
- Lesson Assessment (Quiz)

Network Time Protocol Concepts

Network Time Protocol (NTP) is used to synchronize the time on a network of machines.

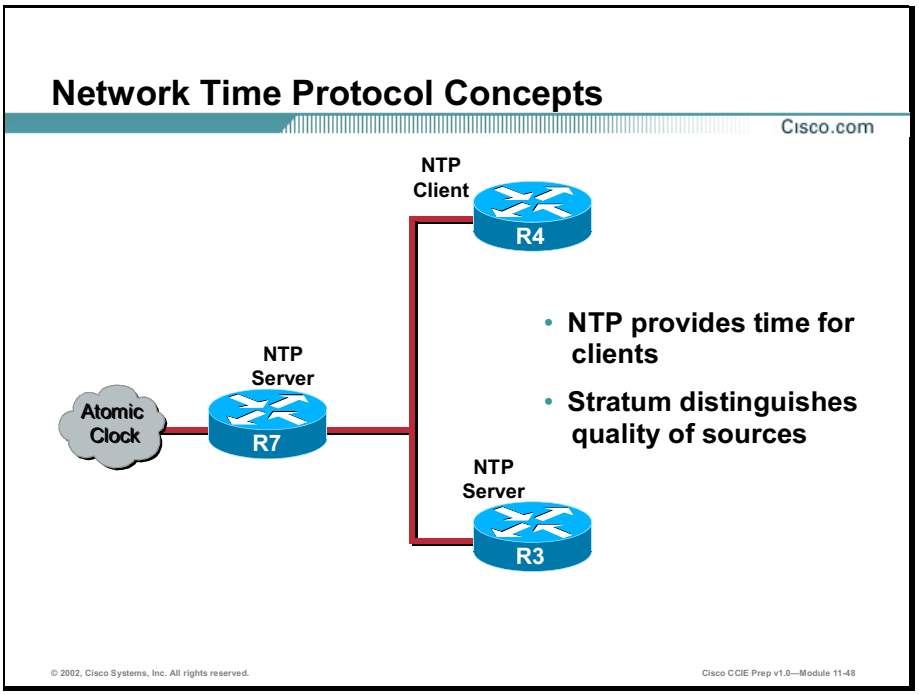

NTP runs over the User Datagram Protocol (UDP) port 123 as both the source and destination. NTP Version 3 is used to synchronize timekeeping among a set of distributed time servers and clients. A set of nodes on a network is identified and configured with NTP and the nodes form a synchronization subnet, sometimes referred to as an overlay network. While multiple masters (primary servers) may exist, there is no requirement for an election protocol.

NTP uses the concept of a "stratum" to describe how many NTP "hops" away a machine is from an authoritative time source. A "stratum 1" time server typically has a radio or atomic clock directly attached, a "stratum 2" time server receives its time via NTP from a "stratum 1" time server, and so on. A machine running NTP will automatically choose the peer with the lowest stratum number as its time source. This strategy effectively builds a self-organizing tree of NTP speakers.

NTP can be setup with devices acting as servers (masters), clients, or peers. A master provides time using a directly attached clocking device (atomic clock, hardware clock, etc.). A client can be configured with the address of the server or listen to the NTP broadcasts on UDP port 123.

Basic Configuration

This section covers the basic configuration of NTP.

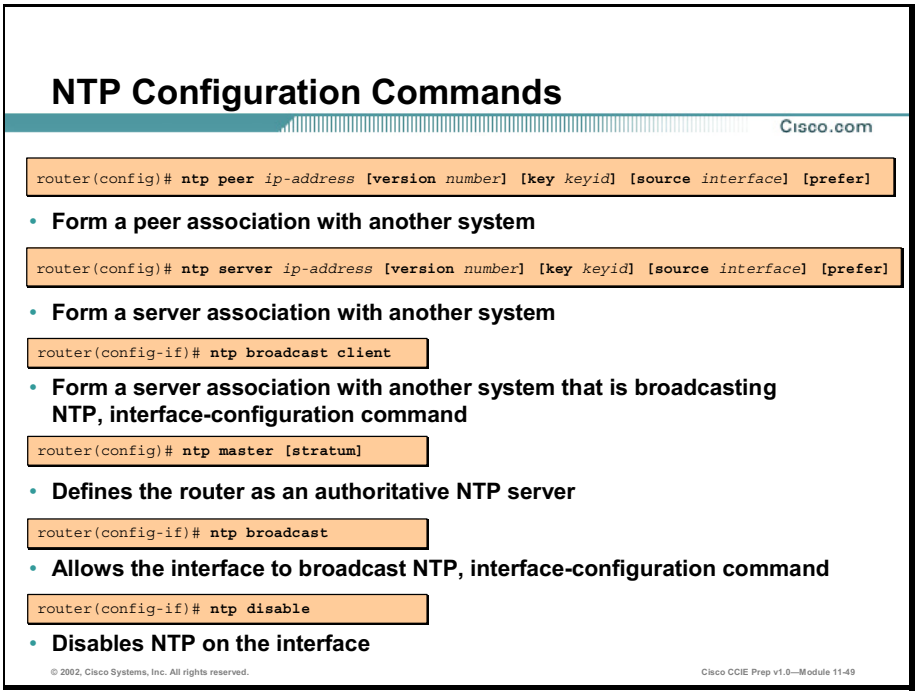

The communications between machines running NTP, known as associations, are usually statically configured; each machine is given the Internet Protocol (IP) address of all machines with which it should form associations. Accurate timekeeping is possible by exchanging NTP messages between each pair of machines with an association. However, in a Local Area Network (LAN) environment, you can configure NTP to use IP broadcast messages. With this alternative, you can configure the machine to send or receive broadcast messages, but the accuracy of timekeeping is marginally reduced because the information flow is one-way only.

Table 10-5: NTP Association Commands

Command	Description
`ntp peer` ***`ip-address`*** `[version` ***`number`***`] [key` ***`keyid`***`] [source` ***`interface`***`] [prefer]`	Form a peer association with another system
`ntp server` ***`ip-address`*** `[version` ***`number`***`] [key` ***`keyid`***`] [source` ***`interface`***`] [prefer]`	Form a server association with another system
`ntp broadcast client`	Form a server association with another system that is broadcasting NTP, interface-configuration command
`ntp master [stratum]`	Defines the router as an authoritative NTP server
`ntp broadcast`	Allows the interface to broadcast NTP, interface-configuration command

Command	Description
`ntp disable`	Disables NTP on the interface

An NTP association can be a peer association (meaning that this system is willing to either synchronize to the other system or to allow the other system to synchronize to it), or it can be a server association (meaning that only this system will synchronize to the other system, and not the other way around).

Only one end of an association needs to be configured; the other system will automatically establish the association.

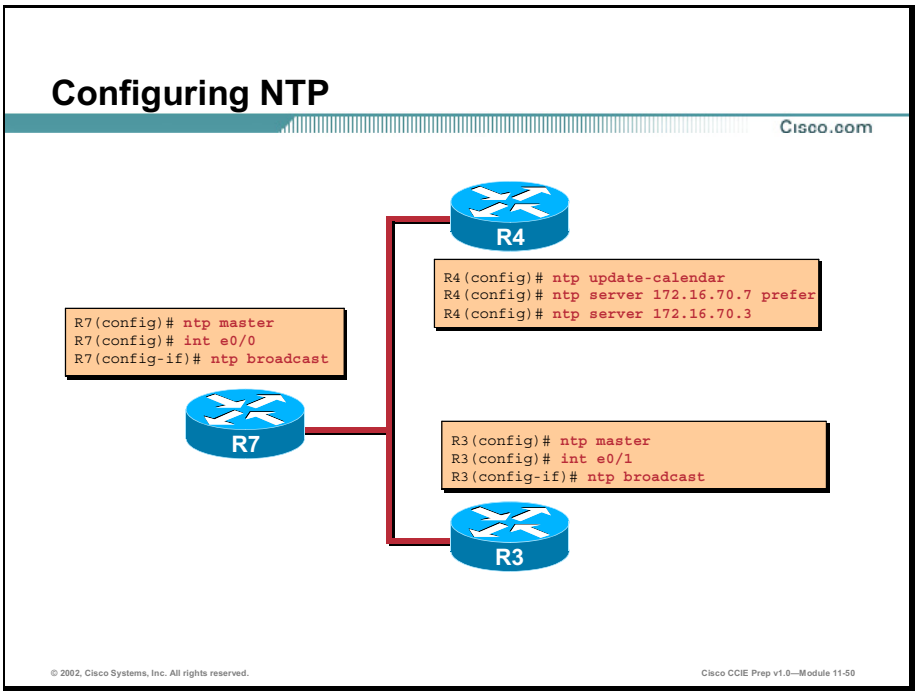

To configure NTP Master, perform the following steps:

Step 1 Define the router that will act as an authoritative NTP server.

```
R7(config)# ntp master
```

Step 2 Verify that the clock is synchronized to the NTP server using **show ntp status**. Inspect the status and time association.

To configure NTP Peer, perform the following steps:

Step 1 Locate an authoritative clock source. For example, you can use a Cisco router or an atomic clock that is attached to a time server. In order to use a Cisco router as the time source, you must manually set that router's local time with the **clock set** command.

Step 2 Configure R4 to use NTP and automatic calendar updates.

```
R4(config)# ntp update-calendar
R4(config)# ntp peer 172.16.70.7 prefer
```

Step 3 Verify that the clock is synchronized to the NTP server using **show ntp status**. Inspect the status and time association.

To configure NTP Broadcast Client, perform the following steps:

Step 1 Locate an authoritative clock source. For example, you can use a Cisco router or an atomic clock that is attached to a time server. Verify the device is sending out NTP broadcasts – the Cisco router interface will have:

```
R7(config)# ntp master
R7(config)# int e0/0
R7(config-if)# ntp broadcast
```

Step 2 Define the interface on the client Router3 to accept NTP Broadcasts and automatic calendar updates.

```
R3(config)# ntp update-calendar
R3(config)# int e0/0
R3(config-if)# ntp broadcast client
```

Step 3 Verify that the clock is synchronized to the NTP server using **show ntp status**. Inspect the status and time association.

To configure a **NTP Server Client**, perform the following steps:

Step 1 Locate an authoritative clock source. For example, you can use a Cisco router or an atomic clock that is attached to a time server.

Step 2 Define the update-calendar on the server Router4 to accept NTP and automatic calendar updates.

```
R4(config)# ntp update-calendar
R4(config)# ntp server 172.16.70.7 prefer
R4(config)# ntp server 172.16.70.3
```

Step 3 Verify that the clock is synchronized to the NTP server using **show ntp status**. Inspect the status and time association.

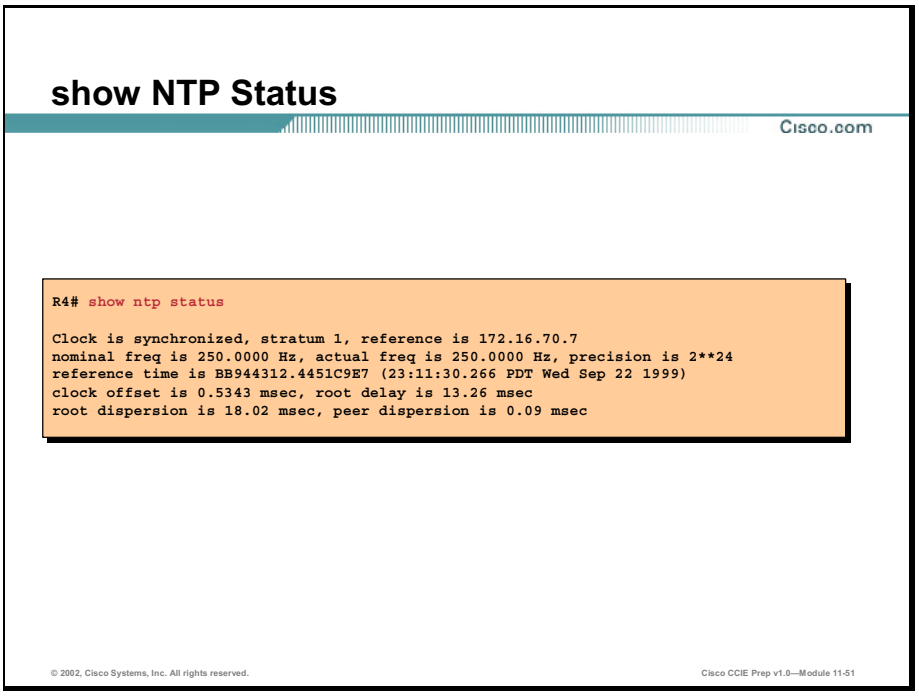

The output indicates this R4 is learning time from the device located at IP address 172.16.70.7 and has a stratum level of 1. Similar output will be found on the other routers.

```
R4# show ntp status

Clock is synchronized, stratum 1, reference is 172.16.70.7
nominal freq is 250.0000 Hz, actual freq is 250.0000 Hz, precision is 2**24
reference time is BB944312.4451C9E7 (23:11:30.266 PDT Wed Sep 22 1999)
clock offset is 0.5343 msec, root delay is 13.26 msec
root dispersion is 18.02 msec, peer dispersion is 0.09 msec
```

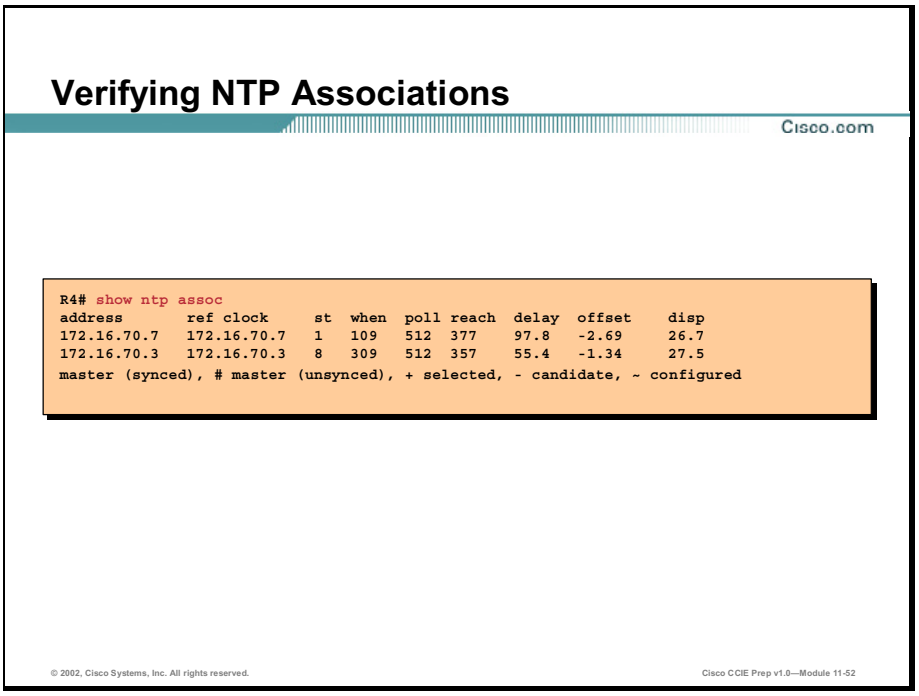

To verify NTP associations, use the **show ntp association** command.

```
R4# show ntp assoc
address         ref clock       st   when   poll reach   delay   offset    disp
172.16.70.7     172.16.70.7     1    109    512  377     97.8    -2.69     26.7
172.16.70.3     172.16.70.3     8    309    512  357     55.4    -1.34     27.5
master (synced), # master (unsynced), + selected, - candidate, ~ configured
```

Authentication Configuration

This section covers the configuration of authentication.

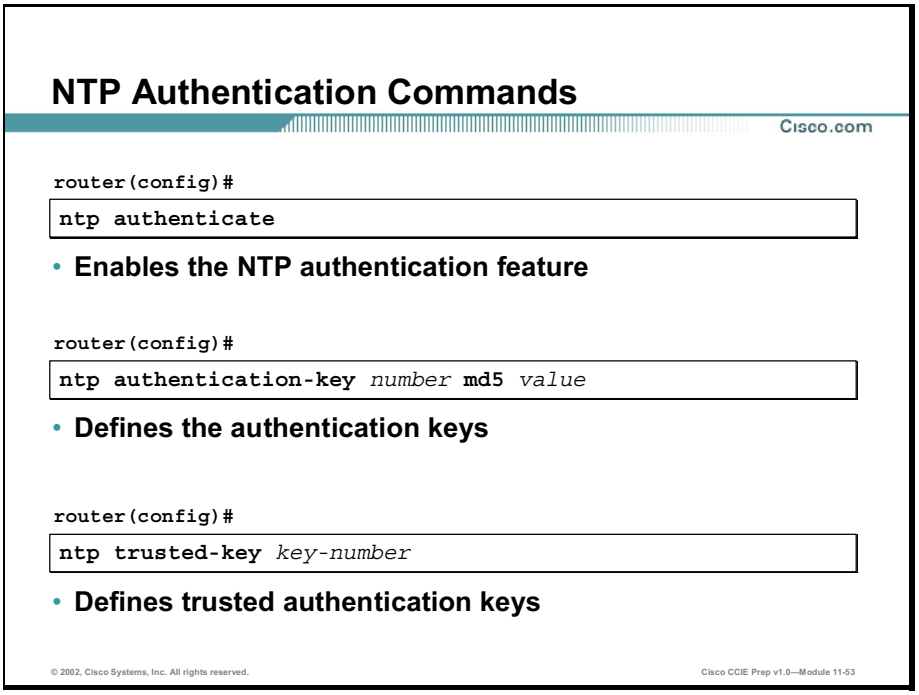

If you want to authenticate the associations with other systems for security purposes, use the commands shown. The first command enables the NTP authentication feature. The second command defines each of the authentication keys. Each key has a key number, a type, and a value. Currently the only key type supported is Message Digest Version 5 (md5). Third, a list of trusted authentication keys is defined. If a key is trusted, this system will be ready to synchronize to a system that uses this key in its NTP packets.

Table 10-6: Authentication Commands

Command	Description
ntp authenticate	Enables the NTP authentication feature
ntp authentication-key *number* md5 *value*	Defines the authentication keys
ntp trusted-key *key-number*	Defines trusted authentication keys

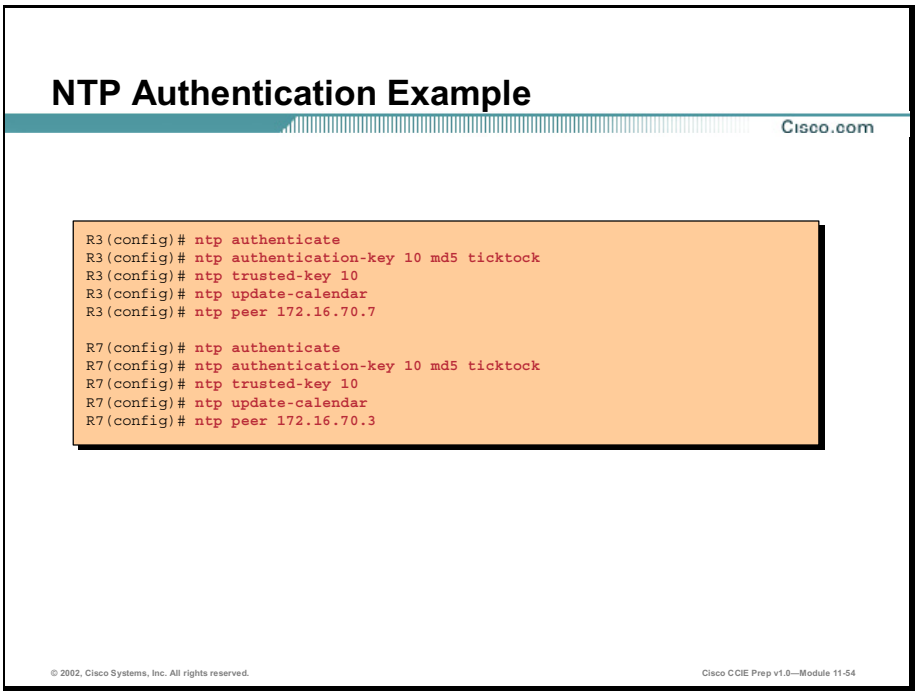

The following example configures routers R3 and R7 as NTP peers.

```
R3(config)# ntp authenticate
R3(config)# ntp authentication-key 10 md5 ticktock
R3(config)# ntp trusted-key 10
R3(config)# ntp update-calendar
R3(config)# ntp peer 172.16.70.7

R7(config)# ntp authenticate
R7(config)# ntp authentication-key 10 md5 ticktock
R7(config)# ntp trusted-key 10
R7(config)# ntp update-calendar
R7(config)# ntp peer 172.16.70.3
```

Timezones

This section covers NTP timezones.

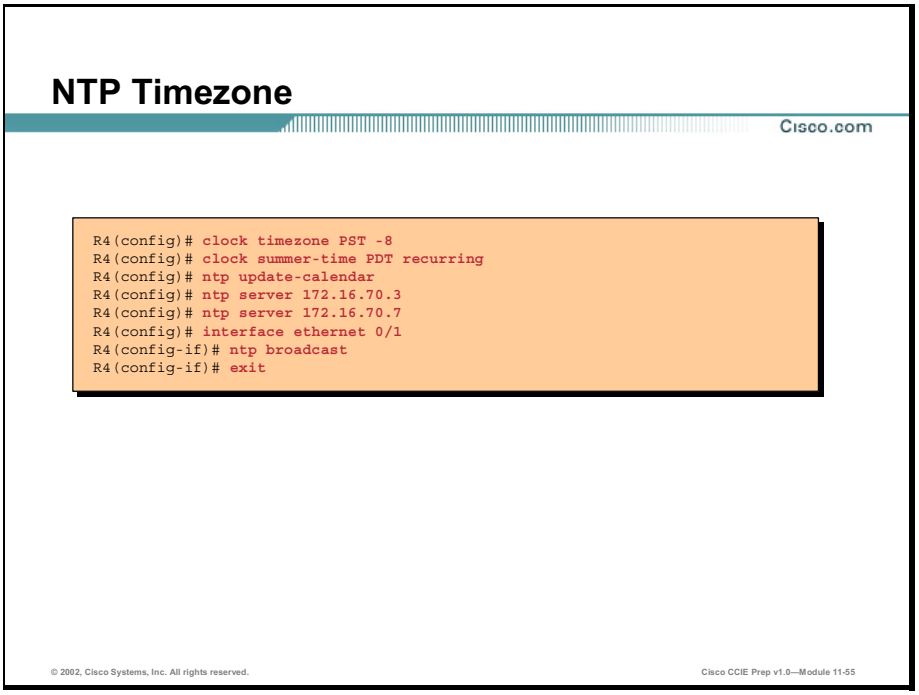

In this example, a router modifies the time from NTP to display with the local timezone offset and supports daylight savings time, periodically updates the calendar, server associations with two other systems, and transmits broadcast NTP packets out interface E1/0.

```
R4(config)# clock timezone PST -8
R4(config)# clock summer-time PDT recurring
R4(config)# ntp update-calendar
R4(config)# ntp server 172.16.70.3
R4(config)# ntp server 172.16.70.7
R4(config)# interface Ethernet 0/1
R4(config-if)# ntp broadcast
R4(config-if)# exit
```

Summary

This section summarizes the key points discussed in this lesson.

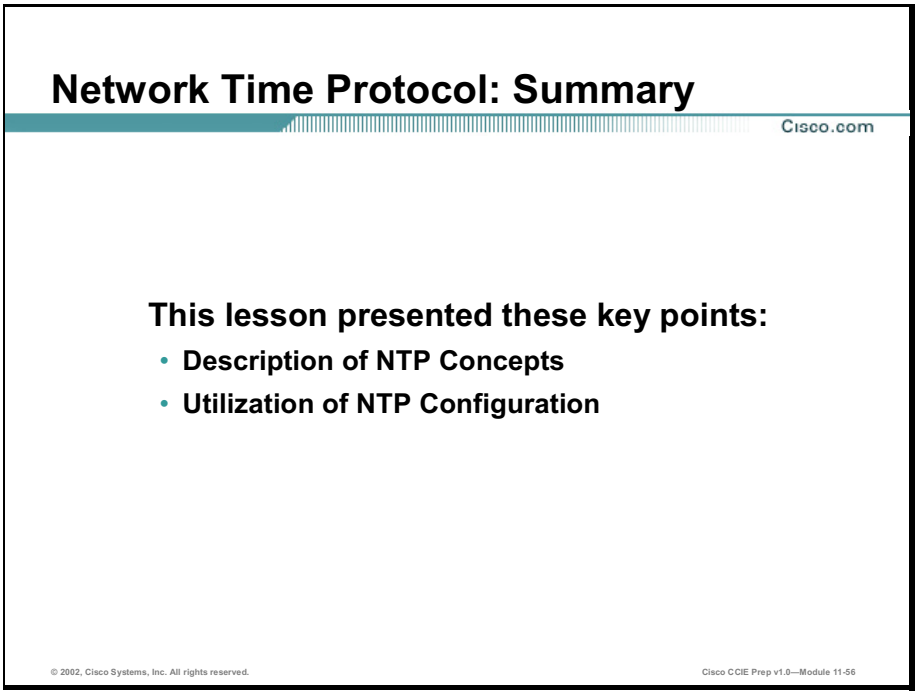

Next Steps

After completing this lesson, go to:

- Network Address Translation

References

For additional information, refer to these resources:

- http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121cgcr/fun_c/fcprt3/fcd303.htm
- RFC 1305

Lesson Assessment (Quiz)

Q1) On receiving multiple NTP broadcasts from a variety of NTP servers at different strata levels, which would the router choose for time synchronization?

A) Stratum 15

B) Stratum 0

C) Stratum 255

D) Stratum 1

E) None of the above

Q2) There are multiple methods to configure NTP on a router. Choose which method is best based on the following information: the router is connected to LAN with 4 NTP servers of different strata levels.

A) router(config)# **ntp server**

B) router(config)# **ntp client**

C) router(config-if)# **ntp broadcast**

D) router(config-if)# **ntp broadcast client**

E) router# **clock set**

F) None of the above, routers are already pre-configured to receive NTP

Q3) Diagnose why NTP authentication is failing between RouterA and RouterB, even though they can ping each other and are directly connected.

RouterA	RouterB

`ntp authenticate`
 	
`ntp authentication-key 10 md5 cisco`
 	
`ntp trusted-key 10`
	`ntp authenticate`
	
	`ntp authentication-key 11 md5 ticktock`
	
	`ntp trusted-key 11`
	
	`ntp peer` **`<Router A IP Address>`**

RouterA	RouterB
ntp peer **<Router B IP Address>**	

A) The trusted key number is wrong on RouterA

B) The md5 value is wrong on Router A

C) Both md5 and trusted-key are wrong on Router A

D) The Routers are not running service timestamp

Q4) A company needs to ensure all routers across the United States provide status in local time, but are synchronized. Which commands should the company use to accomplish this task?

A) Purchase another vendor's routers, for Cisco does not support this function

B) **clock set**

C) **service timestamp**

D) **clock timezone**

E) **clock summer-time**

Q5) True or False: NTP is the only way to set time on a router.

A) True

B) False

Network Address Translation

Overview

Network Address Translation (NAT) allows you to change the source Internet Protocol (IP) address of devices to reduce the need for valid Internet routable addresses.

Importance

Familiarity with NAT is required in order to pass the CCIE Lab Exam.

Objectives

Upon completing this lesson, you will be able to:

- Describe NAT concepts
- Configure NAT
- Verify NAT verification

Learner Skills and Knowledge

To fully benefit from this lesson, you must have these prerequisite skills and knowledge:

- Passed the CCIE written qualification exam
- Completed the Building Cisco Remote Access Networks (BCRAN) course or have the equivalent knowledge

Outline

This lesson includes these sections:

- Overview
- Technology
- Static Translations
- Dynamic Translations
- NAT Overload
- Troubleshooting
- Summary
- Lesson Assessment (Quiz)

Technology

This section covers Network Address Translation (NAT) technology.

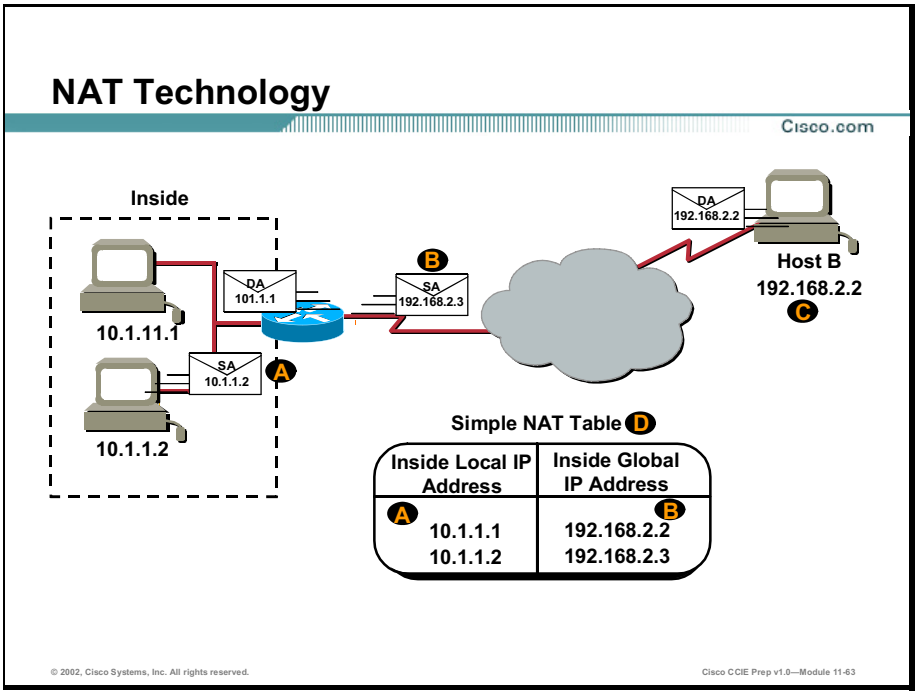

NAT technology enables private IP internetworks that use non-registered IP addresses to connect to a public network such as the Internet. A NAT router is placed on the border of a stub domain (inside network), and a public network (outside network) translates the internal local addresses into globally unique IP addresses before sending packets to the outside network. NAT takes advantage of the fact that relatively few hosts in a stub domain communicate outside of the domain at any given time. Therefore, only a subset of the IP addresses in a stub domain must be translated into globally unique IP addresses for outside communication.

If the internal addresses must change because of changes in service providers or the merger of two intranets (two companies merged, for example), NAT can be used to translate the appropriate addresses. NAT enables address changes dynamically, without changes to hosts or routers other than those bordering stub domains, thereby eliminating duplicate address ranges without readdressing host computers.

The translation performed using NAT can either be static or dynamic. Static translation occurs when you specifically configure addresses in a lookup table. A list of inside addresses maps to a pool of outside addresses. The inside and outside addresses can be statically mapped one-for-one or dynamic mapping can occur. There can be multiple pools of outside addresses. Multiple internal hosts can also share a single outside IP address, which conserves address space. Address sharing is accomplished by port multiplexing, or changing the source port on the outbound packet so that replies can be directed back to the appropriate router.

For load sharing, you can map outside IP addresses to inside IP addresses using the Transmission Control Protocol (TCP) load distribution feature. Load distribution can also be

accomplished using NAT where one external address maps to this address, then the round robin between inside machines occurs. In this case, incoming new connections are distributed across several routers. Each connection may involve information that a given connection must remain on one router.

Cisco's implementation of NAT uses the following terms related to NAT:

Table 10-7: Cisco NAT Terminology

Term	Definition
Inside local IP address (A)	The IP address assigned to a host on the inside network. The address was globally unique but obsolete, allocated from RFC 1918, Address Allocation for Private Internet Space, or randomly picked.
Inside global IP address (B)	A legitimate IP address (assigned by the InterNIC or service provider) that represents one or more inside local IP addresses to the outside world. The address was allocated from a globally unique address space, typically provided by the Internet Service Provider (ISP).
Outside global IP address (C)	The IP address that was assigned to a host on the outside network by its owner. The address was allocated from a globally routable address space.
Outside local IP address	The IP address of an outside host as it appears to the inside network. The address was allocated from address space routable on the inside, or possibly allocated from RFC 1918, for example.
Simple translation (D)	A translation entry that maps one IP address to another.
Extended translation entry	A translation entry that maps one IP address and port pair to another address port pair.

Static Translations

This section lists the steps needed to enable basic, static, and local IP address translation.

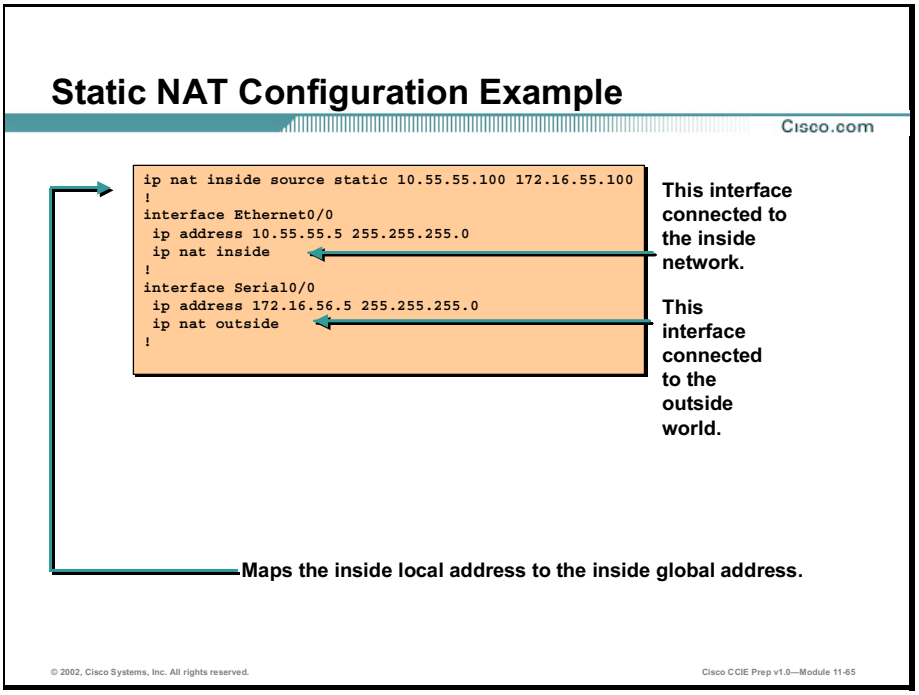

To enable basic, static, and local IP address translation, perform the following steps:

Step 1 At a minimum, IP routing and appropriate IP addresses must be configured on the router.

Step 2 If you are using static address translations for inside local addresses, define the addresses using the **ip nat inside source static** *local-ip global-ip* global configuration command. To remove the static translation, use the **no** form of this command.

Step 3 Define the outside and inside interfaces and apply the configuration appropriately.

Step 4 Enable NAT on at least one inside and one outside interface by entering interface configuration mode and entering the **ip nat** {**inside** | **outside**}command. Only packets moving between inside and outside interfaces can be translated. For example, if a packet is received on an inside interface but is not destined for an outside interface, it will not be translated.

Note In the example above, the NAT pool and the outside interface do not share the same subnet. In order for proper routing to occur, the router(s) directly connected to this interface must be aware of the NAT pool IP addresses for routing purposes.

Table 10-8

Command	Description
local-ip	Sets up a single static translation. This argument establishes the local IP address assigned to a host on the inside network. The address could be randomly chosen, allocated from RFC 1918, or obsolete.
global-ip	Sets up a single static translation. This argument establishes the globally unique IP address of an inside host as it appears to the outside world.

Dynamic Translations

This section lists the steps needed to enable dynamic local IP address translation.

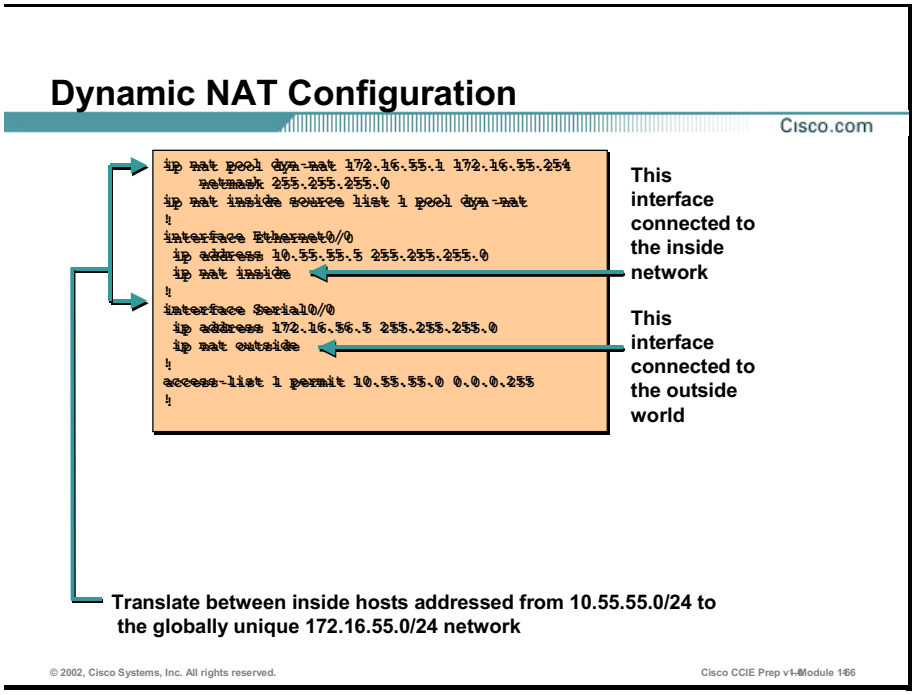

To enable dynamic local IP address translation, perform the following steps:

Step 1 At a minimum, IP routing and appropriate IP addresses must be configured on the router.

Step 2 Define a standard IP access list for the inside network using the **access-list** *access-list-number* {**permit** | **deny**}*local-ip-address* command.

Step 3 Define an IP NAT pool for the inside network using the **ip nat pool** *pool-name start-ip end-ip* {**netmask** *netmask* | **prefix-length** *prefix-length*} [**type rotary**] command.

Step 4 Map the access list to the IP NAT pool using the **ip nat inside source list** *access-list-number* **pool** *pool-name* command.

Step 5 Enable NAT on at least one inside and one outside interface with the **ip nat**{**inside** | **outside**} command.

Step 6 Only packets traveling between inside and outside interfaces can be translated. For example, if a packet is received on an inside interface but is not destined for an outside interface, it will not be translated.

Table 10-9: ip nat pool Commands

Command	Description
pool-name	Name of the pool.

Command	Description
start-ip	Starting IP address that defines the range of addresses in the address pool.
end-ip	Ending IP address that defines the range of addresses in the address pool.
netmask *netmask*	Network mask that indicates which address bits belong to the network and subnetwork fields, and which bits belong to the host field. Specify the netmask of the network to which the address pool belongs.
prefix-length *prefix-length*	Number that indicates how many bits of the netmask are 1s (how many bits of the address indicate the network). Specify the netmask of the network to which the pool addresses belong.
type rotary	(Optional) Indicates that the range of addresses in the address pool identifies real, inside hosts among which TCP load distribution will occur.

NAT Overload

This section describes the steps needed to configure inside global address overloading.

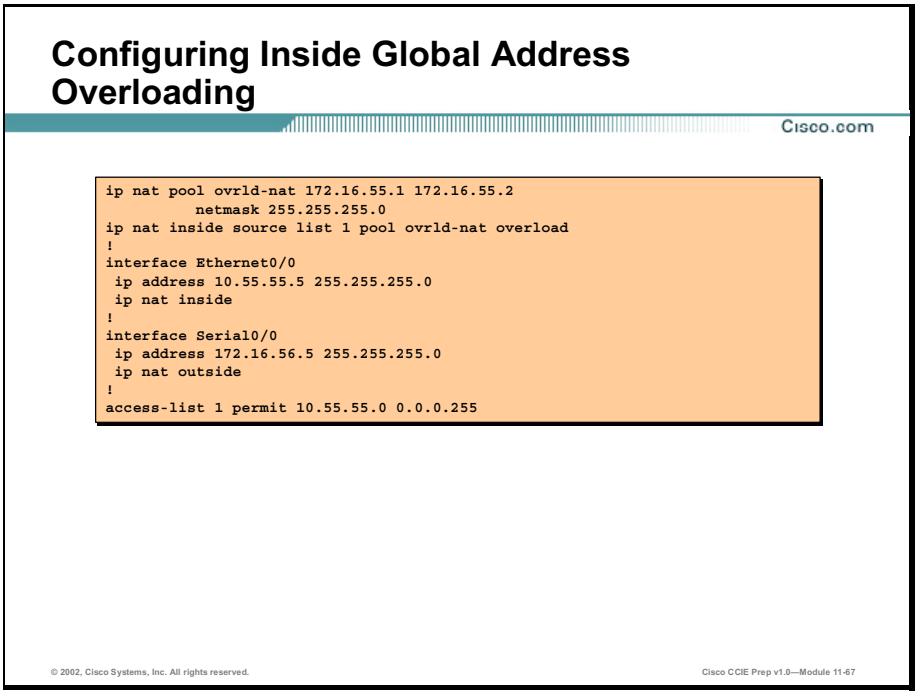

To configure inside global address overloading, perform the following steps:

Step 1 At a minimum, IP routing and appropriate IP addresses must be configured on the router.

Step 2 Configure dynamic address translation.

When you define the mapping between the access list and the IP NAT pool using the **ip nat inside source list** *access-list-numbe*r **pool** *pool-name* **overload** command, add the **overload** keyword to the command.

Step 3 Enable NAT on the appropriate interfaces using the **ip nat** {**inside** | **outside**} command.

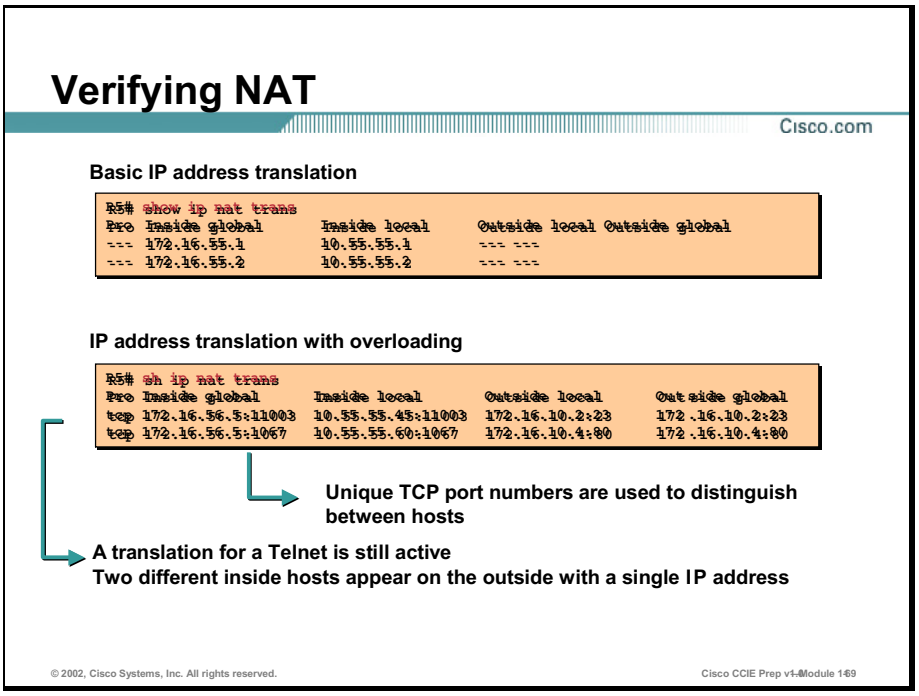

The following commands can be used to verify NAT operation:

Table 10-10

Command	Description
`show ip nat translations [verbose]`	Shows active translations
`show ip nat statistics`	Shows translation statistics

Troubleshooting

This section covers how to troubleshoot NAT.

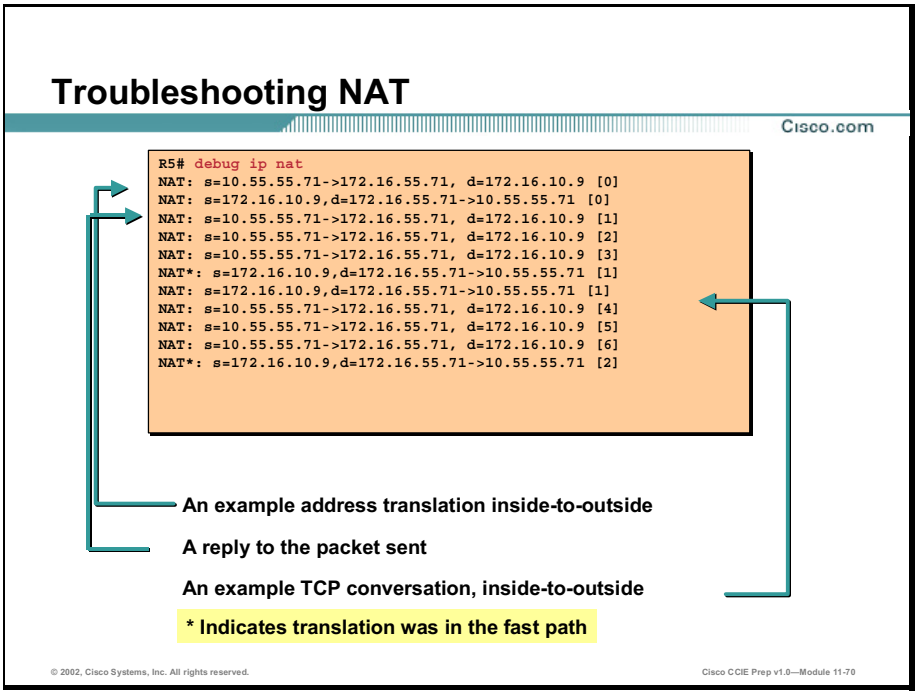

If you need to use a trace on NAT operation, use the following command:

Table 10-11

Command	Description
`debug ip nat [list | detailed]`	Displays a line of output for each packet that gets translated

As shown in the figure, decode the debug output using the following key points:

- The asterisk next to NAT indicates that the translation is occurring in the fast path. The first packet in a conversation will always go through the slow path (be process-switched). The remaining packets will go through the fast path if a cache entry exists
- s=10.55.55.71 is the source address
- d=172.16.10.9 is the destination address
- 10.55.55.71->172.16.55.71 indicates that the address was translated
- The value in brackets is the IP identification number. This information may be useful for debugging because it enables you to correlate with other packet traces from sniffers, for example

Summary

This section summarizes the key points discussed in this lesson.

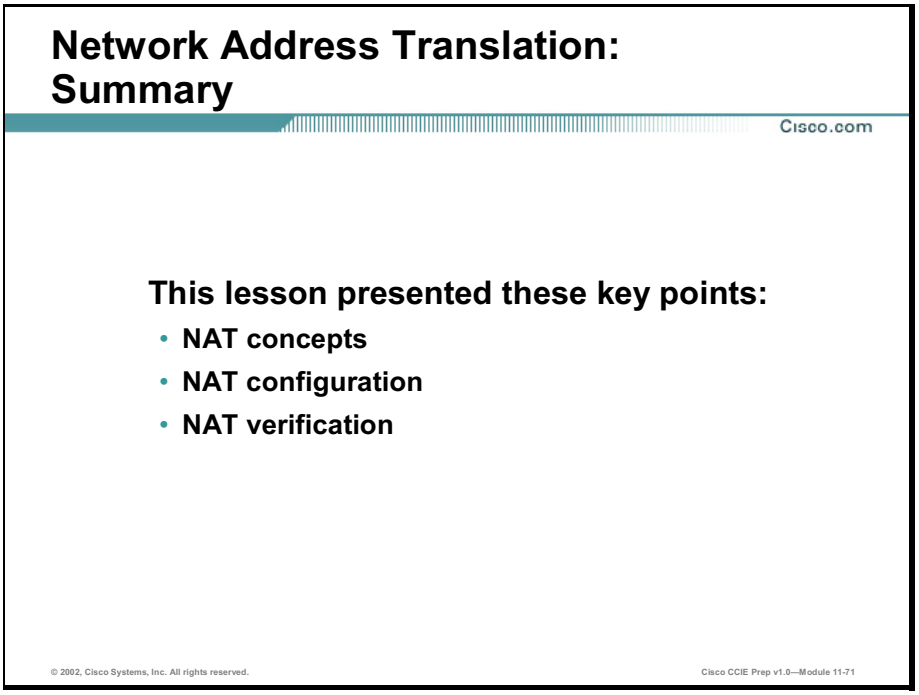

Next Steps

After completing this lesson, go to:

- Hot Standby Routing Protocol

References

For additional information, refer to these resources:

- http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121cgcr/ip_c/ipcprt1/1cdipadr.htm
- RFC 1918
- RFC 1631

Lesson Assessment (Quiz)

Q1) True or False: NAT modifies the destination address IP address.

A) True

B) False

Q2) If the NAT global inside pool is not seen as a group of IP addresses in the outside interface subnet, what action must be taken?

A) A static routing statement on the 'Natting' router must be made

B) The directly connected router to the outside interface must be routing aware of the IP address(es) in the nat global inside pool IP address(es)

C) The configuration must have overload enabled

D) A PIX Firewall should be used

Q3) Upon using a **show ip nat translation** command, there is only one inside global address being utilized. What has happened?

A) Flow Address Translation is enabled

B) Overload is enabled

C) The nat inside global pool is of one address and overload is enabled

D) The access list for the inside local pool is of one address and overload is enabled

Q4) There are 50 inside devices but only 8 valid Internet routable IP addresses. What should the configuration be to allow over 8 inside devices to simultaneously access the Internet?

A) The overload sub-command is required

B) Customer must ask for a larger pool of valid Internet addresses

C) The rotary sub-command is required

D) Use only 8 addresses on the inside devices

E) Define the access list for the inside local pool to only allow 8 devices

Q5) True or False: NAT supports route maps in 12.0 code.

A) True

B) False

Hot Standby Routing Protocol

Overview

Hot Standby Routing Protocol (HSRP) enables two or more routers to appear as a single default gateway for host devices.

Importance

HSRP is required to pass the CCIE Lab.

Objectives

Upon completing this lesson, you will be able to:

- Describe HSRP Concepts
- Utilize HSRP Load Balancing
- Configure HSRP Tracking

Learner Skills and Knowledge

To fully benefit from this lesson, you must have these prerequisite skills and knowledge:

- Passed the CCIE written qualification exam
- Completed the Building Cisco Multilayer Switched Networks (BCMSN) course or have the equivalent knowledge

Outline

This lesson includes these sections:

- Overview
- HSRP Concepts
- HSRP Authentication
- Summary
- Lesson Assessment (Quiz)

HSRP Concepts

This section covers Hot Standby Router Protocol (HSRP) concepts.

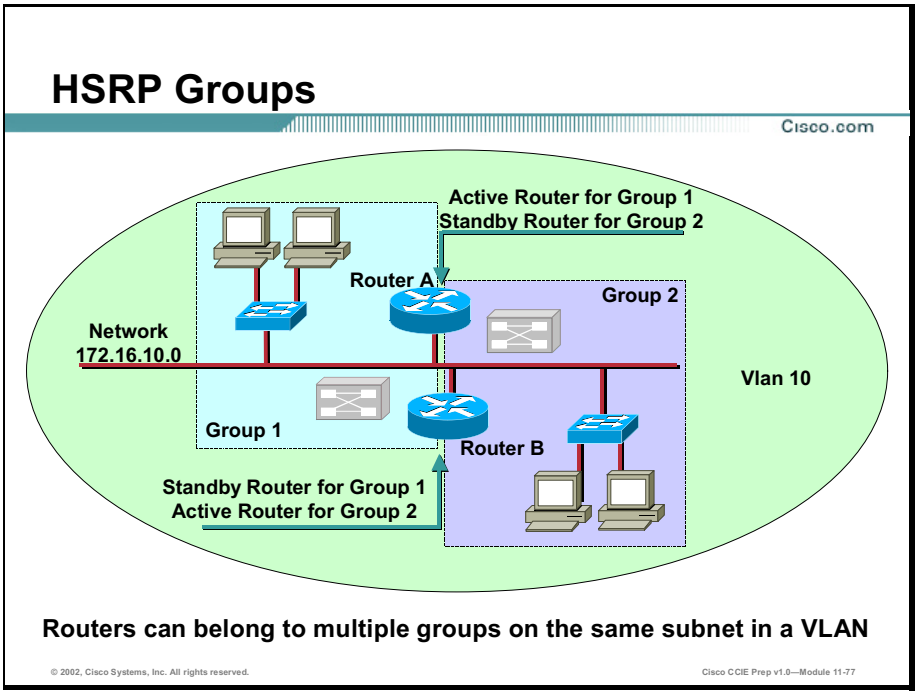

The Hot Standby Router Protocol (HSRP) provides fault tolerance and enhanced routing performance for Internet Protocol (IP) networks. HSRP allows Cisco Internetwork Operating System (IOS) routers to monitor each other's operational status and very quickly assume packet-forwarding responsibility in the event the current forwarding device in the HSRP group fails or is taken down for maintenance. The standby mechanism remains transparent to the attached hosts and can be deployed on any Local Area Network (LAN) type. With multiple hot standby groups, routers can simultaneously provide redundant backup and perform load sharing across different IP subnets

HSRP uses groups to define what routers are participating in handling fault tolerance for device hosts by appearing to the hosts as a single default gateway IP address.

HSRP IP address is configured in all group members and host devices as the default gateway IP address. Further, HSRP will create a Media Access Control (MAC) address for all group members to use on their standby interface.

The HSRP active router can be defined among the HSRP group by priority and can be assured to be the active router as long as it is functional with the preempt command.

Finally, if an interface(s) that provides the host with access to the network off HSRP routers becomes unavailable, an HSRP router can reduce its chances of becoming the active router.

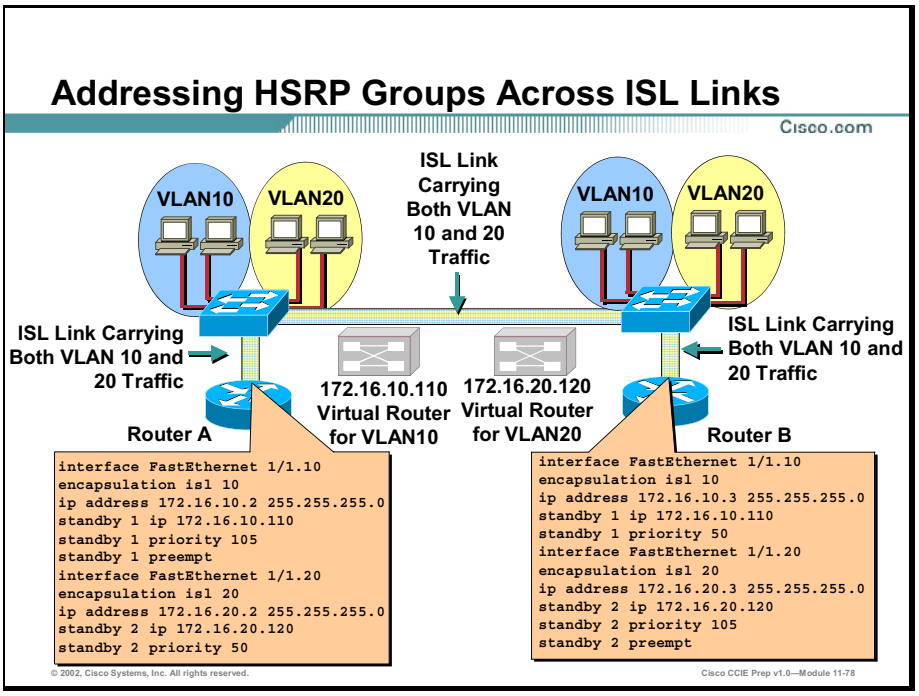

Routers can simultaneously provide redundant backup and perform load sharing across different IP subnets.

For each standby group, an IP address and a single well-known MAC address with a unique group identifier is allocated to the group.

The IP address of a group is in the range of addresses belonging to the subnet in use on the LAN. However, the IP address of the group must differ from the addresses allocated as interface addresses on all routers and hosts on the LAN, including virtual IP addresses assigned to other HSRP groups.

This example shows the configuration for two HSRP-enabled routers participating in two separate virtual LANs (VLANs) using Inter-Switch Link (ISL). Running HSRP over ISL allows users to configure redundancy between multiple routers that are configured as front ends for VLAN IP subnets. By configuring HSRP over ISLs, users can eliminate situations in which a single point of failure causes traffic interruptions. This feature inherently provides some improvement in overall networking resilience by providing load balancing and redundancy capabilities between subnets and VLANs.

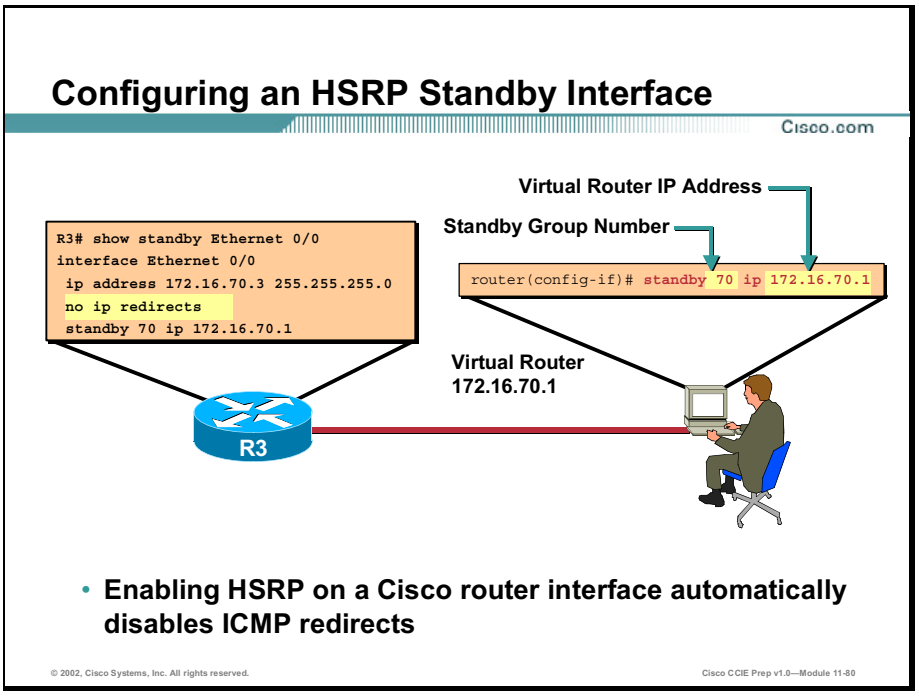

To configure a router as a member of an HSRP standby group, enter the following command in interface configuration mode.

R3(config-if)# standby *group-number* **ip** *ip-address*

Table 10-12

Variable	Definition
group-number	(Optional) Indicates the HSRP group to which this interface belongs. Specifying a unique group number in the standby commands enables the creation of multiple HSRP groups. The default group is 0.
ip-address	Indicates the IP address of the virtual HSRP router.

While running HSRP, it is important that the end user stations do not discover the actual MAC addresses of the routers in the standby group. Any protocol that informs a host of the router actual address must be disabled. To ensure that the actual addresses of the participating HSRP routers are not discovered, enabling HSRP on a Cisco router interface automatically disables Internet Control Message Protocol (ICMP) redirects on that interface.

Once the **standby** command is issued, the interface changes to the appropriate state. When the router successfully executes the command, the router issues an HSRP message. The following is an example of one state message that might be generated.

```
3w1d : %STANDBY-6-STATECHANGE: Standby: 70: Ethernet 0/0 state Speak -> Standby
3w1d : %STANDBY-6-STATECHANGE: Standby: 70: Ethernet 0/0 state Standby -> Active
```

The following example states that interface Ethernet 0/0 is a member of the HSRP standby group 70, the virtual router IP address for that group is 172.16.70.1, and that ICMP redirects are disabled.

```
R3# show run
Building configuration...
Current configuration:
!
(text deleted)
interface Ethernet 0/0
 ip address 172.16.70.3 255.255.255.0
 no ip redirects
 standby 70 ip 172.16.70.1
!
```

To remove an interface from an HSRP group, enter the **no standby** *group* **ip** command.

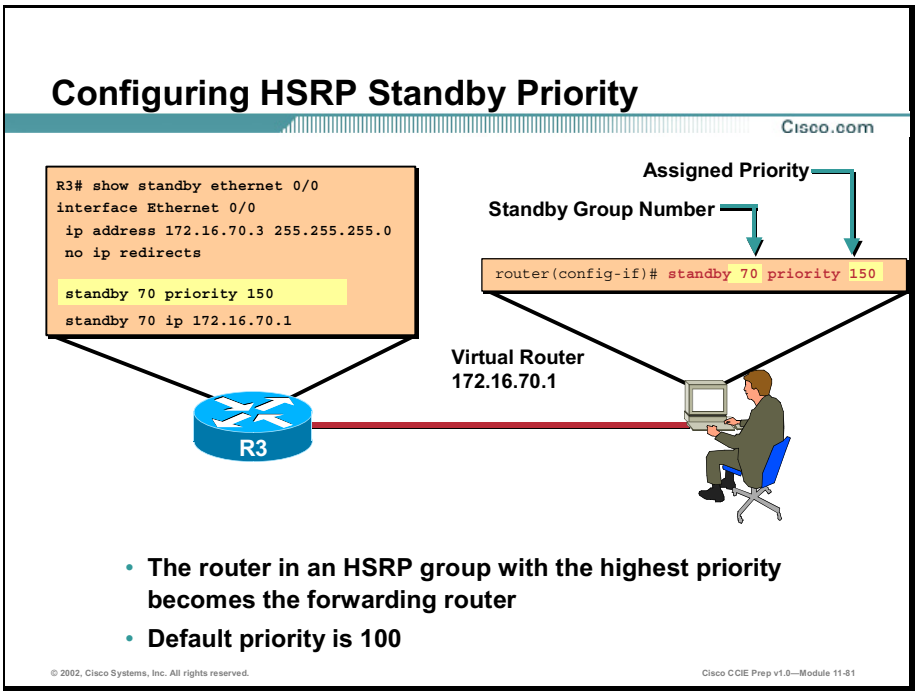

Each standby group has its own active and standby routers. The network administrator can assign a priority value to each router in a standby group, allowing the administrator to control the order in which active routers for that group are selected.

To set the priority value of a router, enter the following command in interface configuration mode.

```
router(config-if)# standby group-number priority priority-value
```

Table 10-13

Variable	Definition
group-number	Indicates the HSRP standby group. This number can be in the range of 0 to 255.
priority-value	Indicates the number that prioritizes a potential Hot Standby router. The range is 0 to 255; the default is 100.

During the election process, the router in an HSRP group with the highest priority becomes the forwarding router. Typically, the active router during configuration is the first router configured for HSRP. If the active and standby routers become unavailable and the remaining routers have the same priority configured, the router with the lowest MAC address becomes the active router. The following example states that interface Ethernet 0/0 has a priority value of 150 in HSRP standby group 70. If this priority value is the highest number in that HSRP standby group, then the routing device on which this interface resides is the active router for that group.

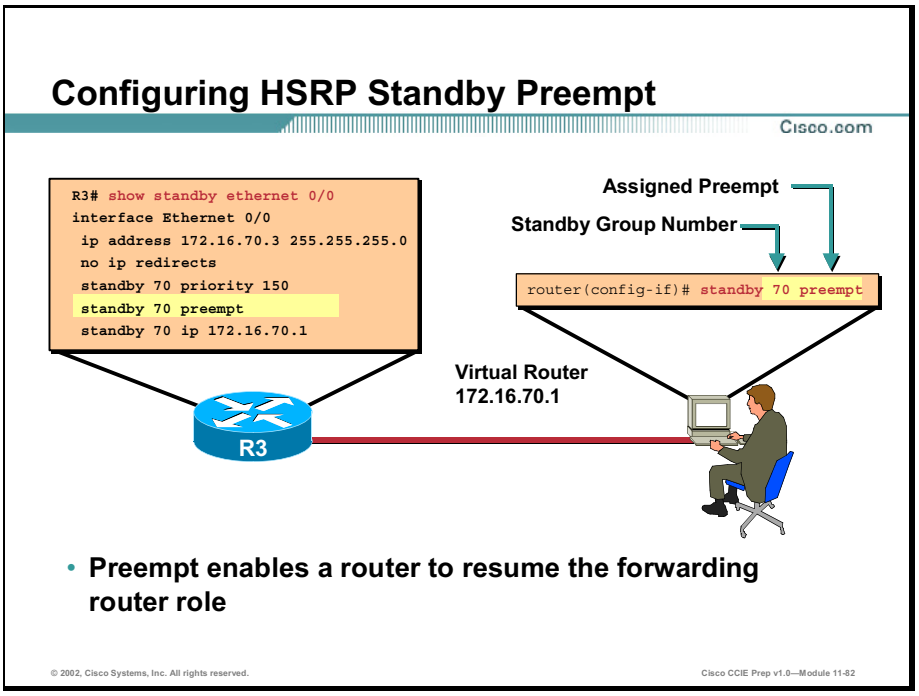

The standby router automatically assumes the active router role when the active router fails or is removed from service. This new active router remains the forwarding router even when the former active router with the higher priority regains service in the network.

The former active router can be configured to resume the forwarding router role from a router with a lower priority. To enable a router to resume the forwarding router role, enter the following command in interface configuration mode.

R3(config-if)# standby *group-number* **preempt**

Once the **standby preempt** command is issued, the interface changes to the appropriate state. The following is an example of the state message generated. This message is automatically generated as soon as the router becomes active in the network.

```
3w1d : %STANDBY-6-STATECHANGE: Standby: 70: Ethernet 0/0 state Standby -> Active
```

The following example states that interface Ethernet 0/0 is configured to resume its role as the active router in HSRP group 70, assuming interface Ethernet 0/0 on this router has the highest priority in that standby group.

```
R3# show run
Building configuration...
Current configuration:
!
(text deleted)
interface Ethernet 0/0
 ip address 172.16.70.3 255.255.255.0
 no ip redirects
 standby 70 priority 150
```

```
standby 70 preempt
standby 70 ip 172.16.70.1
```

To remove the interface from preemptive status, enter the **no standby** *group* **preempt** command.

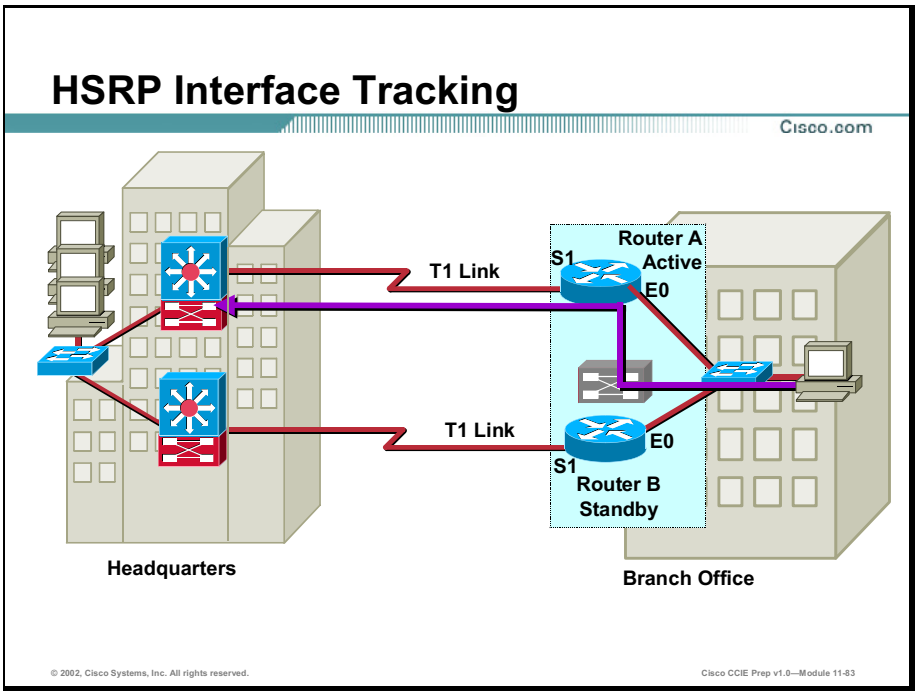

The T1 link between the active forwarding router for the standby group and headquarters experiences a failure. Without HSRP enabled, Router A would detect the failed link and send an ICMP redirect to Router B. However, when HSRP is enabled, ICMP redirects are disabled. Therefore, neither Router A nor the virtual router sends an ICMP redirect and, although the S1 interface on Router A is no longer functional, Router A still communicates hello messages out interface E0 indicating that Router A is still the active router. Packets sent to the virtual router for forwarding to headquarters cannot be routed. Interface tracking enables the priority of a standby group router to be automatically adjusted based on availability of the interfaces of that router. When a tracked interface becomes unavailable, the HSRP priority of the router is decreased. The HSRP tracking feature reduces the likelihood that a router with an unavailable key interface will remain the active router. In this campus LAN example, the E0 interface on Router A tracks the S1 interface. If the link between the S1 interface and headquarters fails, the router automatically decrements its priority on that interface and stops transmitting hello messages out interface E0. Router B assumes the active router role when no hello messages are detected for the specific holdtime period.

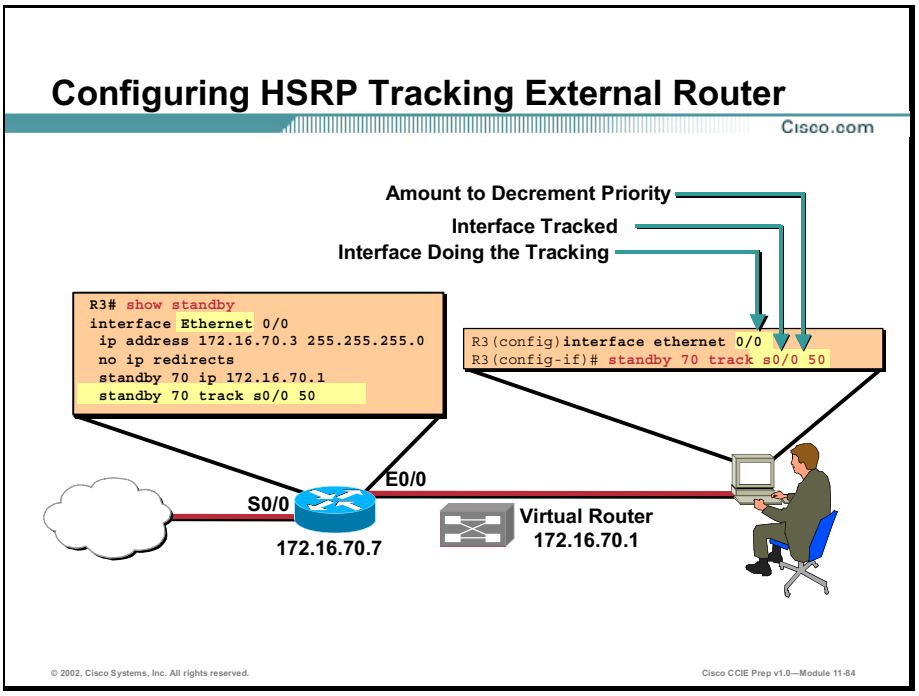

To configure HSRP tracking, enter the following command in interface configuration mode.

R3(config-if)# standby *group-number* **track** *type number interface-priority*

Table 10-14

Variable	Description
group-number	(Optional) Indicates the group number on the interface to which the tracking applies. The default number is 0.
type	Indicates the interface type (combined with the interface number) that will be tracked.
number	Indicates the interface number (combined with the interfaceType) that will be tracked.
interface-priority	(Optional) Indicates the amount by which the hot standby priority for the router is decremented when the interface becomes disabled. The priority of the router is incremented by this amount when the interface becomes available. **The default value is 10**.

To disable interface tracking, enter the **no standby** *group* **track** command.

HSRP Authentication

This section describes the configuration of HSRP authentication.

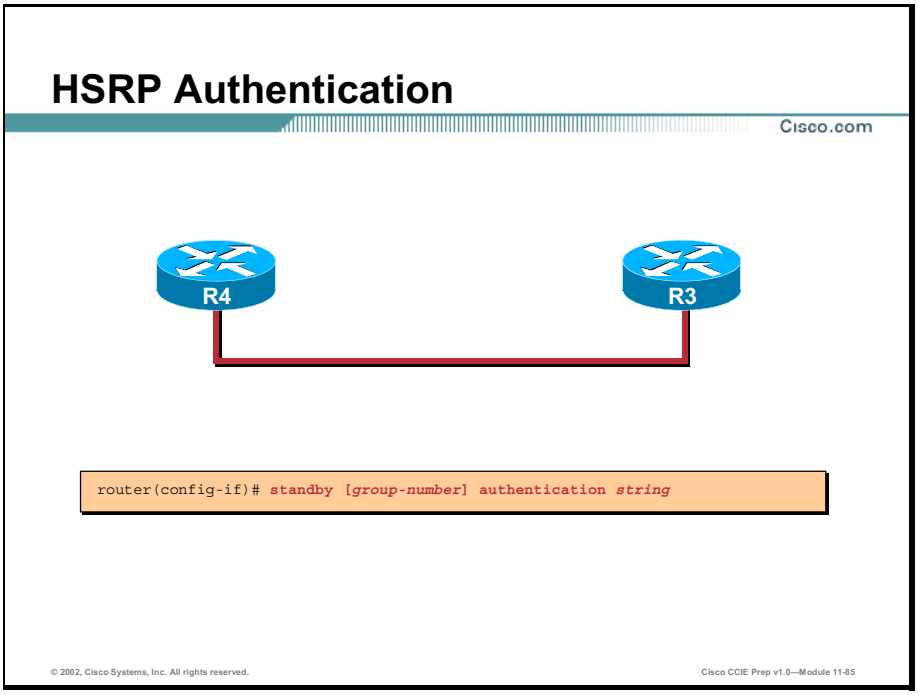

R3(config-if)# standby [*group-number*] authentication *string*

The HSRP authentication feature consists of a shared clear-text key contained within the HSRP packets. The purpose of this password is to disallow mis-configured routers from participating in an HSRP group it was not intended to participate in.

To configure the HSRP authentication string, use the **standby [*group-number*] authentication *string*** command.

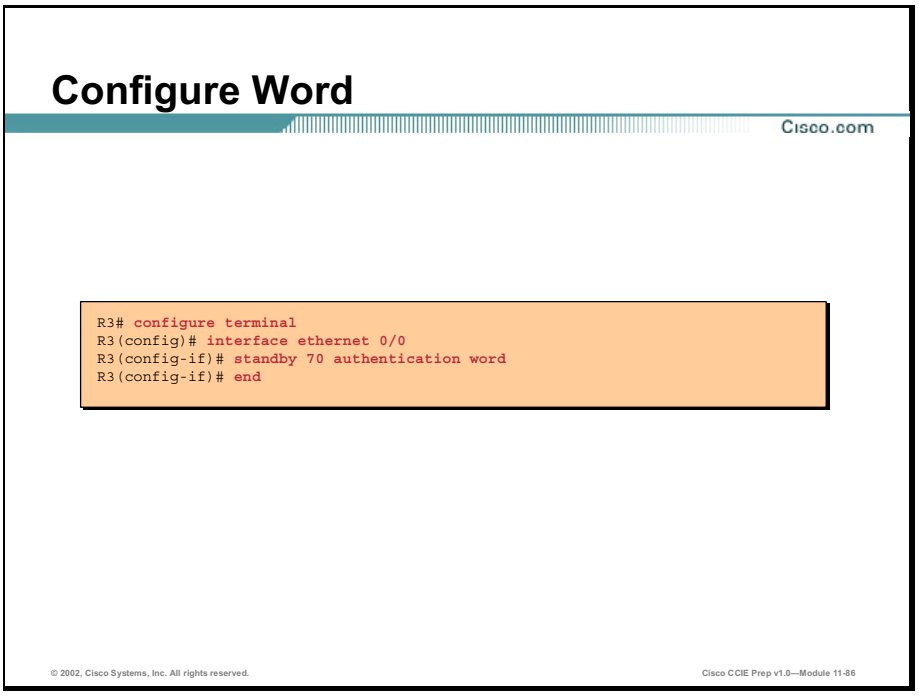


```
R3# configure terminal
R3(config)# interface ethernet0/1
R3(config-if)# standby 70 authentication word
R3(config-if)# end
```

This example shows how to configure *word* as the authentication string required to allow Hot Standby routers in group 70 to interoperate.

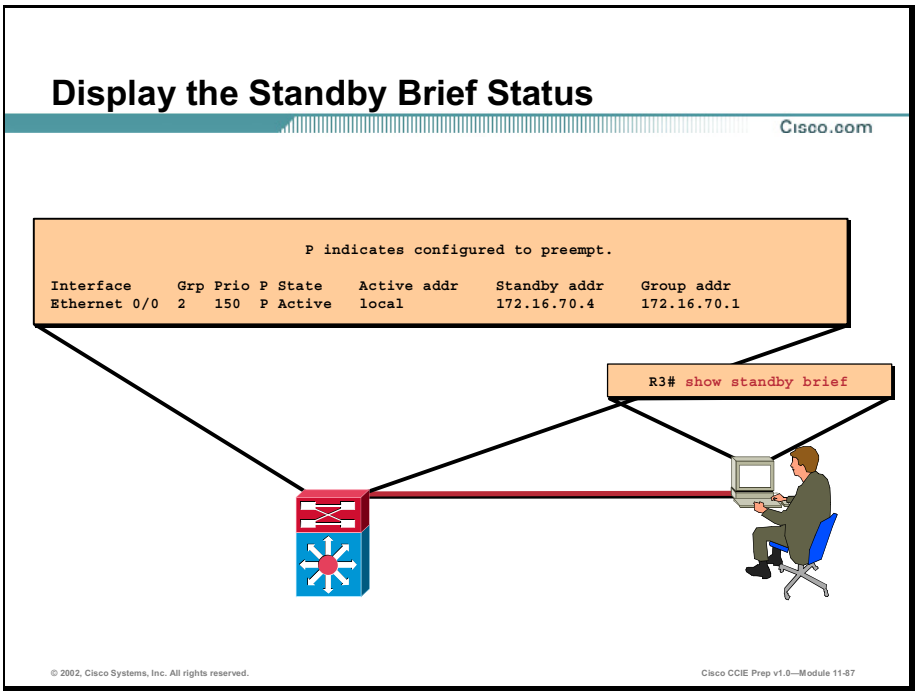

To display the status of the HSRP router, enter the following command in privileged EXEC mode.

R3# show standby *type-number group* **brief**

Table 10-15

Variable	Description
type-number	(Optional) Indicates the target interface type and number for which output is displayed
Group	(Optional) Indicates a specific HSRP group on the interface for which output is displayed
brief	(Optional) Displays a single line of output summarizing each standby group

If the above optional interface parameters are not indicated, the **show standby** command displays HSRP information for all interfaces. Below is an example of the output that results when you specify the *type-number and group* parameters.

```
R3# show standby Ethernet 70
Ethernet 0/0 - Group 70
Local state is Active, priority 150, may preempt
Hellotime 3 holdtime 10
Next hello sent in 00:00:02.944
Hot standby IP address is 172.16.70.1 configured
Active router is local
Standby router is 172.16.70.4 expires in 00:00:08
Standby virtual mac address is 0000.0c07.ac46
```



```
Tracking interface states for 1 interface, 1 up:
Up Serial 0/0 Priority decrement: 40
```

Below is an example of the output resulting when you specify the **brief** parameter.

```
R3# show standby brief
Interface    Grp Prio P State  Active addr Standby addr Group addr
Ethernet 0/0 70  150  P Active local       172.16.70.4  172.16.70.1
```

Note	When specifying a group, you must designate an interface.

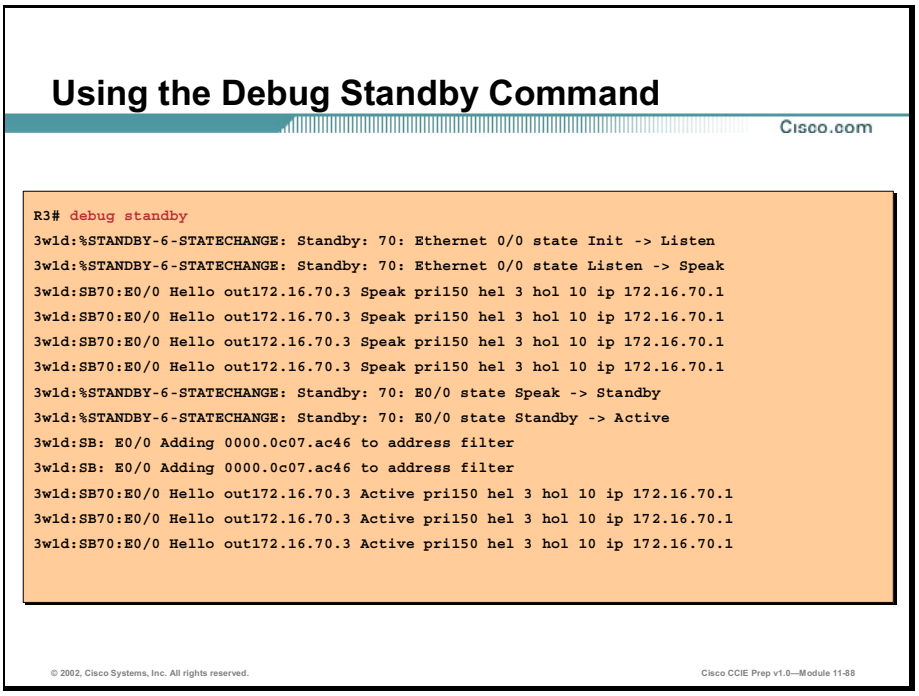

The Cisco IOS implementation of HSRP supports the **debug** command. Enabling the debug facility displays the HSRP state changes and debugging information regarding transmission and receipt of HSRP packets. To enable HSRP debugging, enter the following command in privileged EXEC mode.

```
R3# debug standby
```

Caution	Because debugging output is assigned high priority in the CPU process, this command can render the system unusable.

The following example displays the **debug standby** command output as the router with the IP address 172.16.70.82 initializes and negotiates for the role of the active router.

```
R3# debug standby
3w1d:%STANDBY-6-STATECHANGE: Standby: 70: Ethernet 0/0 state Init -> Listen
3w1d:%STANDBY-6-STATECHANGE: Standby: 70: Ethernet 0/0 state Listen -> Speak
3w1d:SB70:E0/0 Hello out172.16.70.3 Speak pri150 hel 3 hol 10 ip 172.16.70.1
3w1d:SB70:E0/0 Hello out172.16.70.3 Speak pri150 hel 3 hol 10 ip 172.16.70.1
3w1d:SB70:E0/0 Hello out172.16.70.3 Speak pri150 hel 3 hol 10 ip 172.16.70.1
3w1d:SB70:E0/0 Hello out172.16.70.3 Speak pri150 hel 3 hol 10 ip 172.16.70.1
3w1d:%STANDBY-6-STATECHANGE: Standby: 70: E0/0 state Speak -> Standby
3w1d:%STANDBY-6-STATECHANGE: Standby: 70: E0/0 state Standby -> Active
3w1d:SB: E0/0 Adding 0000.0c07.ac46 to address filter
3w1d:SB: E0/0 Adding 0000.0c07.ac46 to address filter
3w1d:SB70:E0/0 Hello out172.16.70.3 Active pri150 hel 3 hol 10 ip 172.16.70.1
3w1d:SB70:E0/0 Hello out172.16.70.3 Active pri150 hel 3 hol 10 ip 172.16.70.1
```

To disable the debugging feature, enter either the **no debug standby** or **the no debug all** command.

Summary

This section summarizes the key points discussed in this lesson.

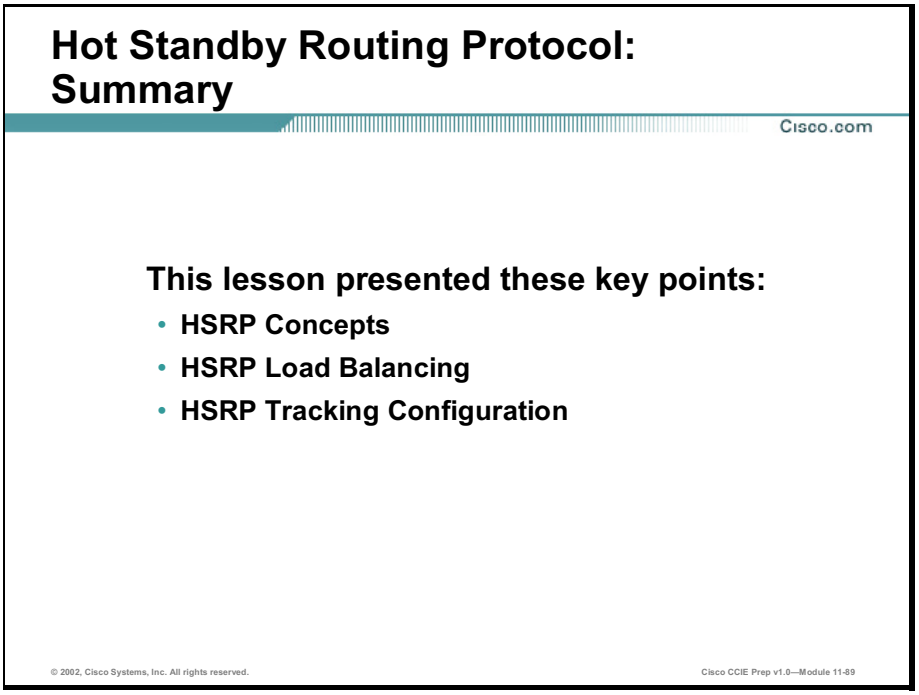

Next Steps

After completing this lesson, go to:

- Dynamic Host Configuration Protocol

References

For additional information, refer to these resources:

- http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121cgcr/ip_c/ipcprt1/1cdip.htm
- *Building Cisco Multi-layered Switched Networks* by Karen Webb

Lesson Assessment (Quiz)

Q1) If four routers are being configured for the same group and have the default stand by priority, which device becomes the active HSRP router?

A) The router with the highest MAC

B) The router with the lowest MAC

C) The first router configured

D) The router with highest IP address

Q2) To ensure that the router for HSRP Group 44 with the highest priority will be the active router, which command must be added to the configuration?

A) router(config-if)# **standby 44 preempt**

B) router(config-if)# **standby 44 ip**

C) router(config-if)# **standby 44 track**

D) router(config-if)# **standby 44 authenticate**

E) router(config-if)# **standby 44 active**

Q3) To configure HSRP load balancing, between two ISL capable routers, _____.

A) Each router supports the other HSRP group, with each router being the active router for one group

B) Overload is enabled

C) The routing protocol will handle load balancing with maximum paths statement

D) The hosts are divided evenly and point to the appropriate router IP address

Q4) HSRP Tracking provides which features?

A) Performs load balancing

B) Allows hosts to track the HSRP multicast

C) Ensures the active router is available

D) Reduces the likelihood that a router with an unavailable key interface will remain the active router

Q5) True or False: HSRP is a routing protocol.

A) True

B) False

Dynamic Host Configuration Protocol

Overview

Dynamic Host Configuration Protocol (DHCP) uses a client/server architecture to provide an Internet Protocol (IP) address to an unconfigured host.

Importance

DHCP enables the management of IP devices, and is associated to the CCIE Lab.

Objectives

Upon completing this lesson, you will be able to:

- Describe DHCP concepts
- Configure DHCP with options
- Troubleshoot DHCP

Learner Skills and Knowledge

To fully benefit from this lesson, you must have these prerequisite skills and knowledge:

- Passed the CCIE written qualification exam
- Completed the Interconnecting Cisco Network Devices (ICND) course or have the equivalent knowledge

Outline

This lesson includes these sections:

- Overview
- DHCP Concepts
- DHCP Commands
- DHCP Service
- DHCP Show and Debug
- Summary
- Lesson Assessment (Quiz)

DHCP Concepts

This section describes basic Dynamic Host Configuration Protocol (DHCP) concepts.

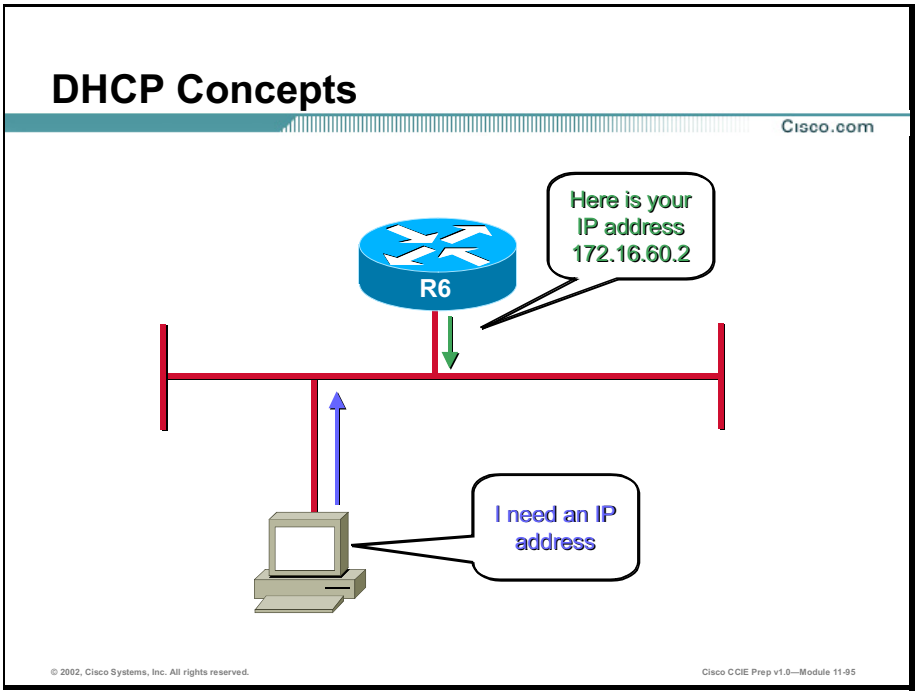

Dynamic Host Configuration Protocol (DHCP) enables you to automatically assign reusable Internet Protocol (IP) addresses to DHCP clients. The Cisco Internetwork Operating System (IOS) DHCP Server feature is a full DHCP server implementation that assigns and manages IP addresses from specified address pools within the router to DHCP clients. If the Cisco IOS DHCP Server cannot satisfy a DHCP request from its own database, it can forward the request to one or more secondary DHCP servers defined by the network administrator.

DHCP requires that a pool of IP addresses (known as the DHCP Scope) be defined as those handed out by the server. This pool can have exclusions within it, such as statically configured IP addresses. Typically, DHCP is defined for a directly connected network for which the router is the default gateway. In addition to the host IP address and subnet mask, DHCP Options are provided by the DHCP Server to define the Default Gateway, Domain Name System (DNS) Servers, Windows Internet Naming Service (WINS) Servers, and lease-time for the IP address.

DHCP Commands

This section describes configuration of the DHCP address pool subnet and mask.

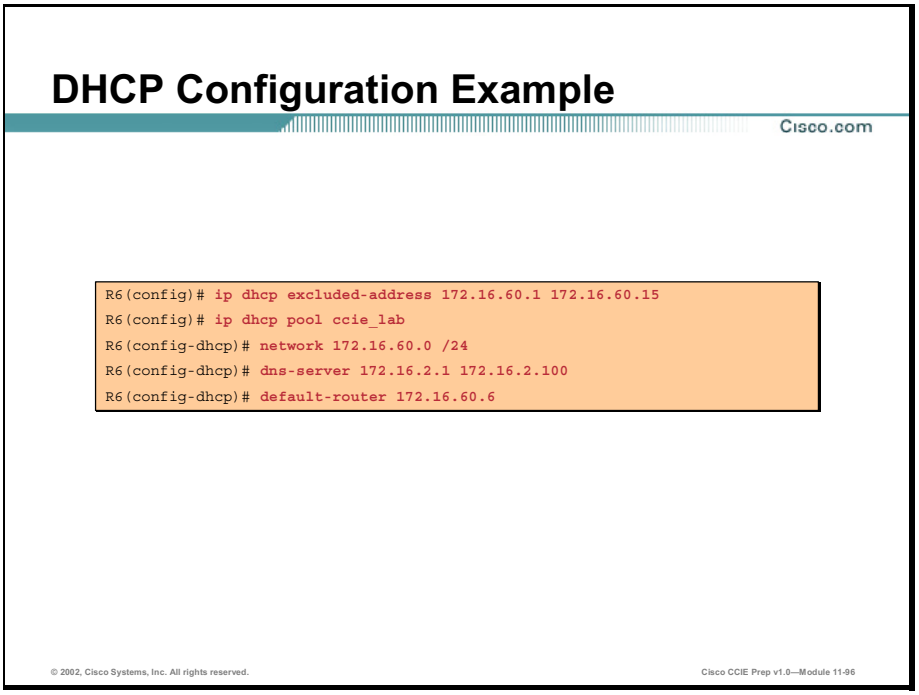

The above configuration will set up a DHCP pool called ccie_lab. The network range will be the 172.16.60.0 network. Addresses 172.16.60.1 – 172.16.60.15 will be excluded from the pool. The following additional parameters will also be configured: default gateway will be set to 172.16.60.6 and DNS servers will be 172.16.2.1 and 172.16.2.100.

The following table lists the DHCP commands and a description of the function of each:

Table 10-16: DHCP Commands

Command	Description
`R6(config)# ip dhcp pool name`	Creates a name for the DHCP server address pool and places you in DHCP pool configuration mode (identified by the config-dhcp# prompt).
`R6(config-dhcp)# network network-number [mask | /prefix-length]`	Specifies the subnet network number and mask of the DHCP address pool. The prefix length specifies the number of bits that comprise the address prefix. The prefix is an alternative way of specifying the network mask of the client. The prefix length must be preceded by a forward slash (/).
`R6(config-dhcp)# default-router ip-address`	Species the default router the DHCP client will use.

Command	Description
R6(config-dhcp)# **dns-server** *address* [*address2 ... address8*]	Specifies the IP address of a DNS server that is available to a DHCP client. One IP address is required; however, you can specify up to eight IP addresses in one command line.
R6(config-dhcp)# **netbios-name-server** *address* [*address2 ... address8*]	Specifies the NetBIOS WINS server that is available to a Microsoft DHCP client. One address is required; however, you can specify up to eight addresses in one command line.
R6(config-dhcp)# **lease** {*days* [*hours*][*minutes*] | **infinite**}	Specifies the duration of the lease. The default is a one-day lease.
R6(config)# **ip dhcp excluded-address** *low address* [*high address*]	Specifies the IP addresses that DHCP will not provide.

DHCP Service

This section describes enabling the Cisco IOS DHCP Server on your router.

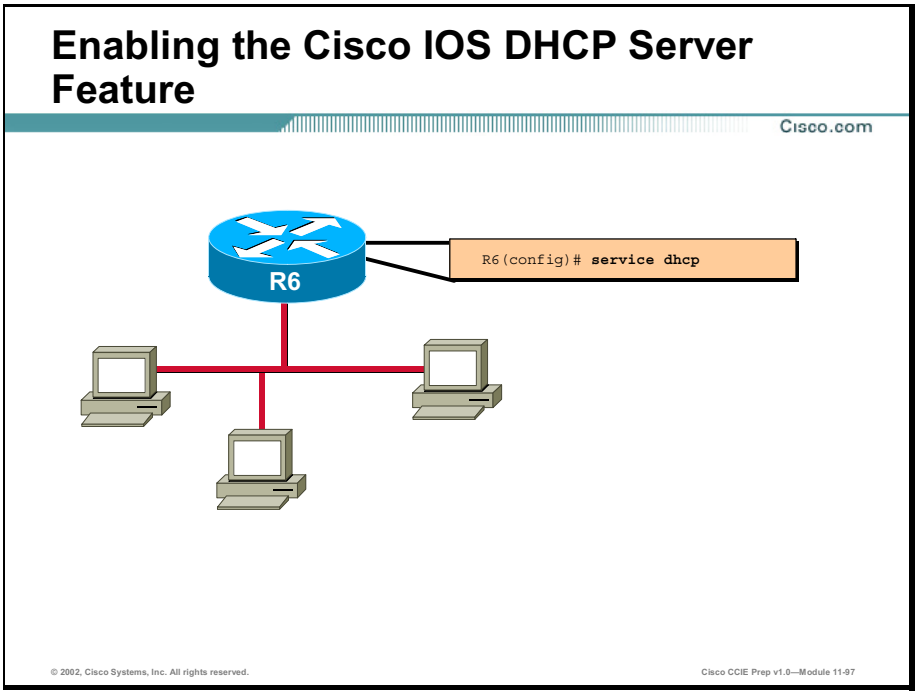

By default, the Cisco IOS DHCP Server feature is enabled on your router. If the feature is disabled, use the following command in global configuration mode to re-enable the Cisco IOS DHCP Server feature on your router:

Table 10-17: < router(config)# service dhcp> Command

Command	Description
`R6(config)# service dhcp`	Enables the Cisco IOS DHCP Server feature on your router. Use the **no** form of this command to disable the Cisco IOS DHCP Server feature.

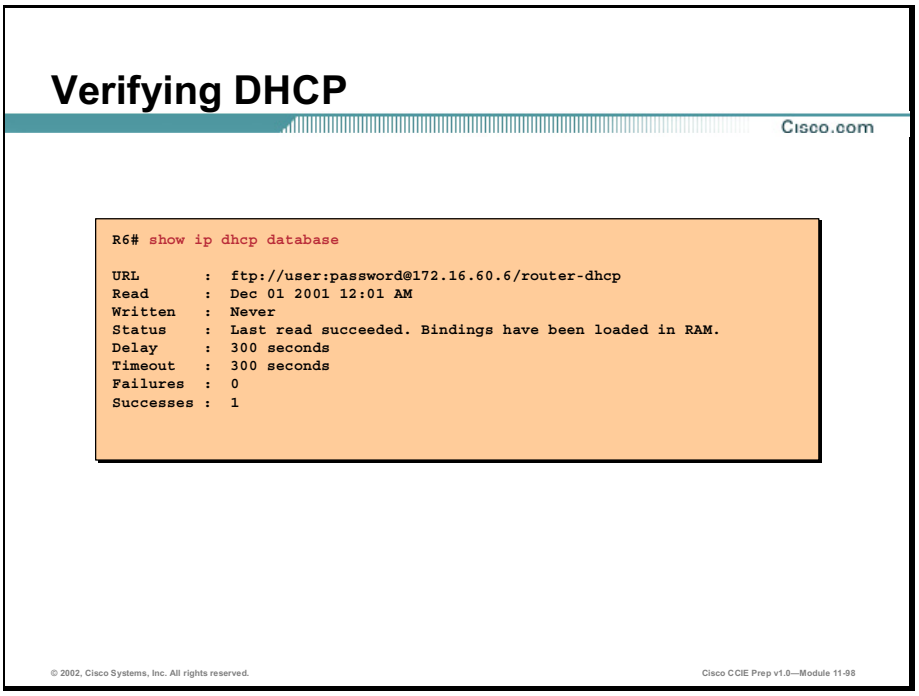

You can verify DHCP using the following commands:

Table 10-18 <show ip dhcp> Command

Command	Description
R6# **show ip dhcp database** [*url*]	Displays recent activity on the DHCP database. **Note:** Use this command in privileged EXEC mode.
R6# **show ip dhcp server statistics**	Displays count information about server statistics and messages sent and received.

DHCP Show and Debug

To enable DHCP server debugging, use the **debug ip dhcp server** command in privileged EXEC mode.

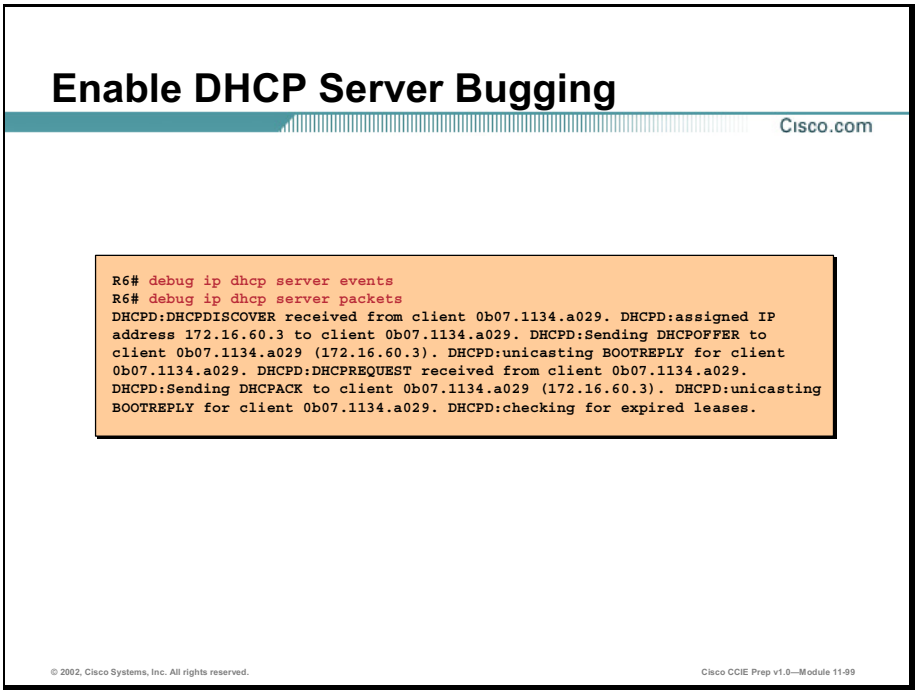

The following command enables debugging on the DHCP server:

Table 10-19: `R6# debug ip dhcp server {events | packets | linkage}>` **Command**

Command	Description
`R6# debug ip dhcp server {events | packets | linkage}`	Enables debugging on the DHCP server.

Summary

This section summarizes the key points discussed in this lesson.

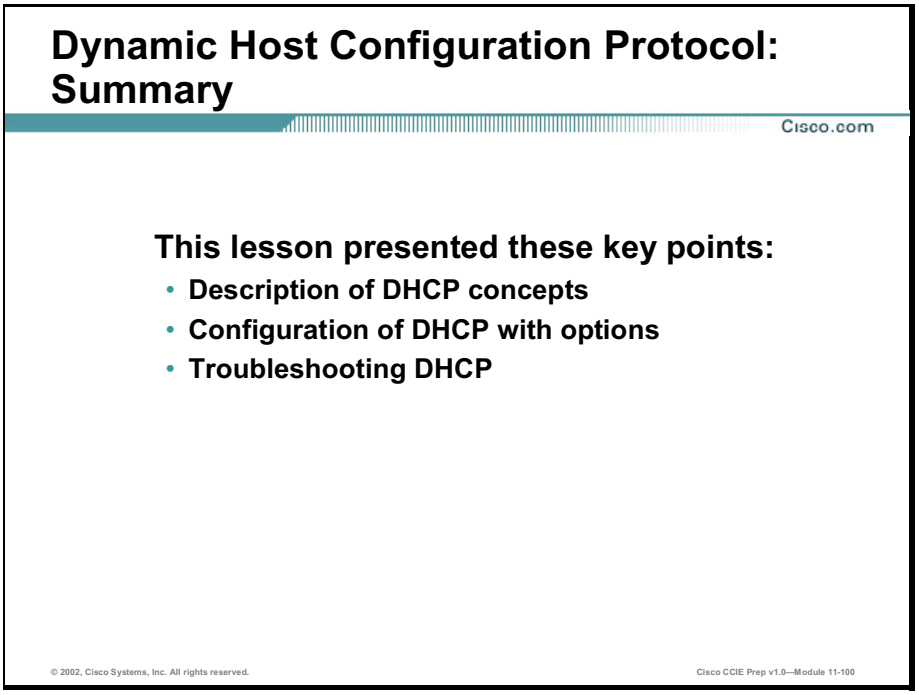

Next Steps

After completing this lesson, go to:

- Security, VoIP, and QoS

References

For additional information, refer to these resources:

- http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121cgcr/ip_c/ipcprt1/1cddhcp.htm
- RFC 2131

Lesson Assessment (Quiz)

Q1) Which DHCP command enables DHCP on the router?

A) router(config)# **ip dhcp pool**

B) router(config)# **network**

C) router(config)# **service ip dhcp**

D) router(config)# **service dhcp**

Q2) Why should the default router IP address be excluded from the DHCP pool?

A) The DHCP Server does not inadvertently create an address conflict

B) The default gateway is a unique device that receives DHCP from a separate mechanism

C) It provides additional IP addresses for NAT

D) The default-router IP address is used as an option

Q3) DHCP is used to aid in management of _____.

A) Login and passwords

B) Router IP address ranges

C) IP addresses

D) Logins

11

Security, VoIP, and QoS

Overview

This module covers Security, Voice over IP (VoIP), and Quality of Service (QoS) as it relates to the CCIE Lab. These are not core topics, but are usually laid on top of the core network infrastructure. You should have a thorough understanding of these topics before attempting the CCIE lab.

Upon completing this module, you will be able to:

- Secure a network using the Security features available in the Cisco IOS
- Configure Voice over IP between two Cisco routers
- Implement Quality of Service features on Cisco routers

Outline

The module contains these lessons:

- Security Concepts
- Voice over IP Concepts
- Quality of Service Concepts

Security Concepts

Overview

This lesson discusses security topics related to the CCIE lab exam. Topics covered include; controlling access to the router, traffic filtering, and prevention of Denial of Service (DoS) attacks. This lesson also introduces advanced topics such as Context-Based Access Control (CBAC) and IP Security Encryption (IPSec).

Importance

Security tasks are normally laid on top of the network in the CCIE lab exam. Therefore, they are usually isolated tasks in the lab. However, incorrectly implementing a security feature may disrupt network connectivity, affecting other areas of the lab.

Objectives

Upon completing this lesson, you will be able to:

- Control access to a Cisco router
- Configure privilege levels
- Evaluate privilege level configurations
- Describe Access Control Lists
- Describe Context-Based Access Control (CBAC)
- Create an IPSec Tunnel between two routers

Learner Skills and Knowledge

To fully benefit from this lesson, you must have these prerequisite skills and knowledge:

- Passed the CCIE written qualification exam
- Completed the Managing Cisco Network Security (MCNS) course or have the equivalent knowledge

Outline

This lesson includes these sections:

- Overview
- Controlling access to a Cisco router
- Configuring privilege levels
- Privilege level configuration examples
- Access Control Lists
- Context-Based Access Control (CBAC)
- IPSec
- Summary
- Lesson Assessment (Quiz)

Controlling Access to a Cisco Router

This section discusses the use of usernames and passwords to control access to a Cisco router.

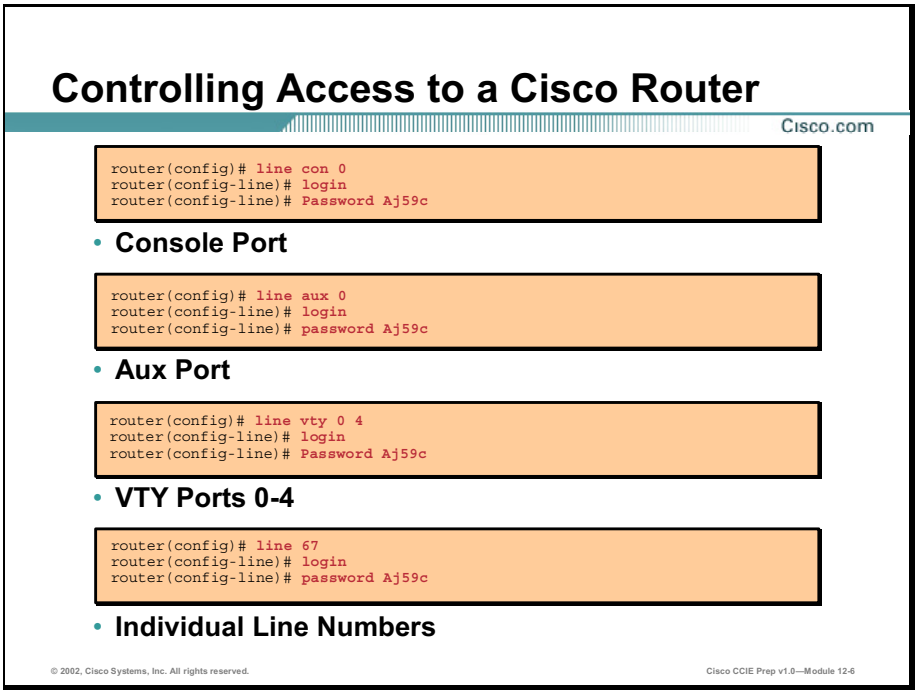

Character (exec) mode on a router can be accessed in the following ways: console port, aux port, Virtual Terminal (VTY) lines using Telnet or Secure Shell (SSH), and Teletype (TTY) lines using a modem. To protect initial access to user mode on the router, you can assign a password to these lines. Use the following commands to access line configuration mode for the appropriate port.

- Console port – **line con 0**
- Aux port – **line aux 0**
- VTY lines – **line vty 0 4**
- TTY lines – **line <line number>**

Line numbers can be determined from the output of the **show line** command. To set or change a password on a line, use the following command in line configuration mode:

Command	Description
`router(config-line)#` **`login`**	Required to set up a shell for the user to enter a password.
`router(config-line)#` **`password`** *`password`*	Set the password required to access the router via this line.

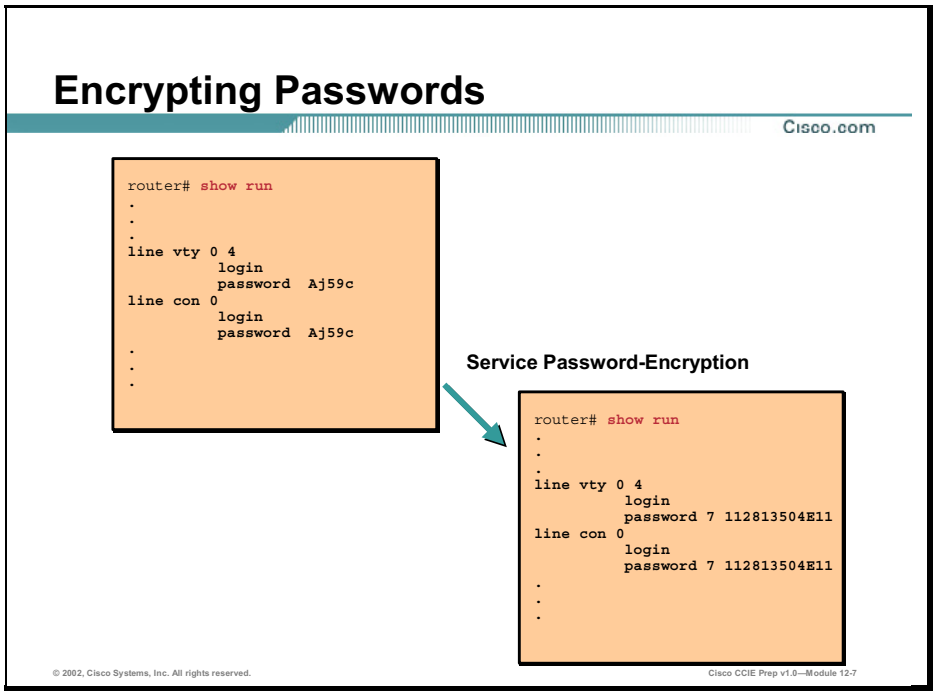

Line passwords and the enable password appear in the router's configuration file in clear-text. To prevent onlookers from reading the passwords, use the **service password-encryption** command. This command will encrypt those passwords so they are viewed in a level 7 encrypted format, rather than clear-text.

Password encryption is applied to all unencrypted passwords in the router's configuration file, including authentication key passwords, the enable password, console, aux, TTY and VTY line access passwords, Point-to-Point (PPP) authentication passwords, and Border Gateway Protocol (BGP) neighbor passwords. Once you have encrypted a password you cannot unencrypt it.

To configure the Cisco Internetwork Operating System (IOS) software to encrypt passwords, use the following command in global configuration mode:

Table 12-2: < service password-encryption > Command

Command	Description
`service password-encryption`	Encrypts all clear-text passwords in the router's configuration file

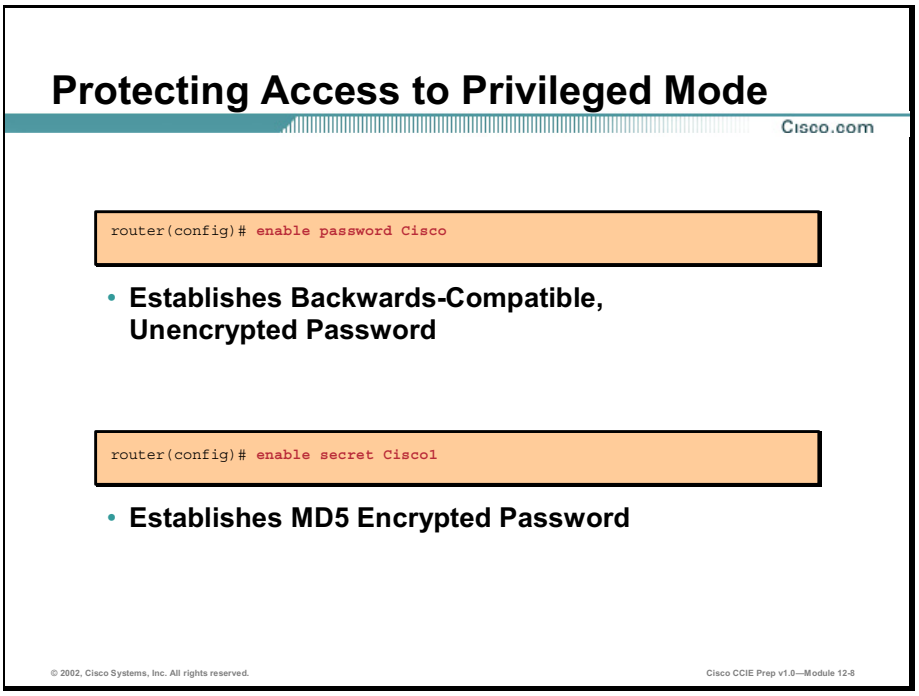

Access to privileged mode on a router can be protected with either the **enable password** or **enable secret** commands. It is recommended that you use the **enable secret** command because it uses an improved Message Digest Version 5 (MD5) encryption algorithm. Use the **enable password** command only if you have an older image of the Cisco Internetwork Operating System (IOS) software, or older boot Read-Only Memory (ROMs) that do not recognize the **enable secret** command.

To configure the router to require a password to access privileged mode, use either of the following commands in global configuration mode:

Table 12-3: Access Privileged Mode Commands

Command	Description
`router(config)# enable password [level privilege level] {password | encryption-type encrypted-password}`	Establishes an unencrypted enable password to access privileged mode
`router(config)# enable secret [level privilege level] {password | encryption-type encrypted-password}`	Specifies an enable secret password, saved using a non-reversible MD5 encryption method

Both of these commands support the **level** keyword to define a password for a specific privilege level. If you specify a password and assign a privilege level, that password can then be used to access that privilege level. Privilege levels are normally used to give special users a subset of privileged exec commands to perform their jobs. Privilege levels will be covered later in this lesson.

If you specify an encryption type, you must provide the password in its encrypted form, meaning an encrypted password you copied from another router's configuration.

If you configure the **enable password** and **enable secret**, the **enable secret** takes precedence over the **enable password** command. These two commands cannot be in effect simultaneously.

Note	You cannot recover a lost enable secret password. You must bypass the configuration file in NonVolatile Random-Access Memory (NVRAM) on the router and set a new enable secret password. See appendix E for password recovery procedures.

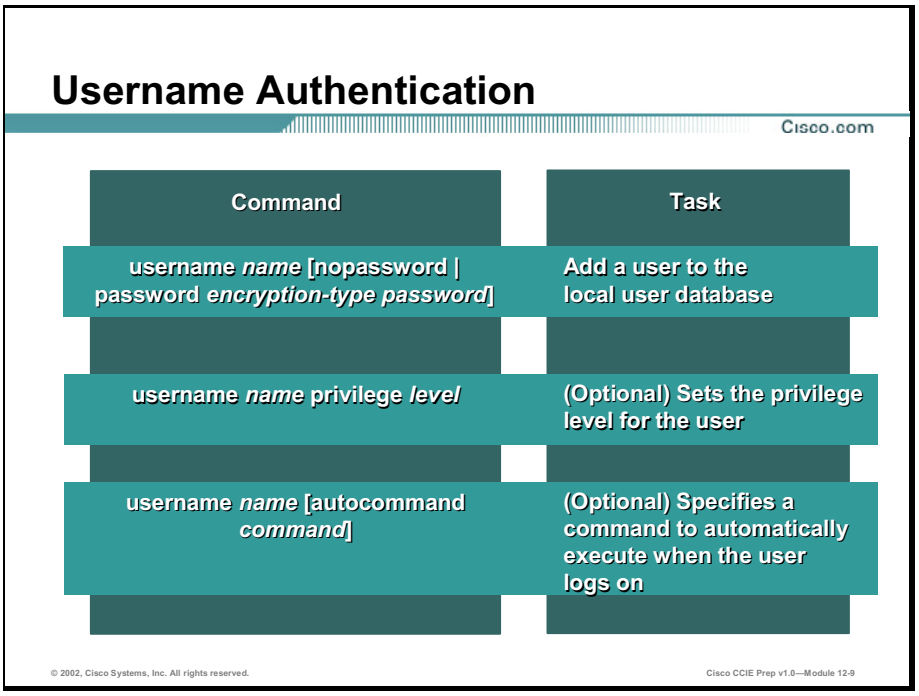

To add an additional layer of security, you can require users to enter a username and password to access the router. This username and password is used in addition to any line or enable passwords in place.

In order for the router to prompt users for a username and password, the following items are required:

- The command **login local** must be entered under line configuration mode for lines that will require a username and password.
- A user must be defined in the router's local user database.

To enable user-based authentication, perform the following steps:

Table 12-4: Enable User-Based Authentication Commands

Step	Command	Description
Step 1	**username** *name* [**nopassword** | **password** *encryption-type* **password**]	Adds a user to the local user database
Step 2	**username** *name* **privilege** *level*	(Optional) Sets the privilege level for the user
Step 3	**username** *name* [**autocommand** *command*]	(Optional) Specifies a command to automatically execute when the user logs on

Configuring Privilege Levels

This section examines the use of privilege levels to customize a user's access level to the router.

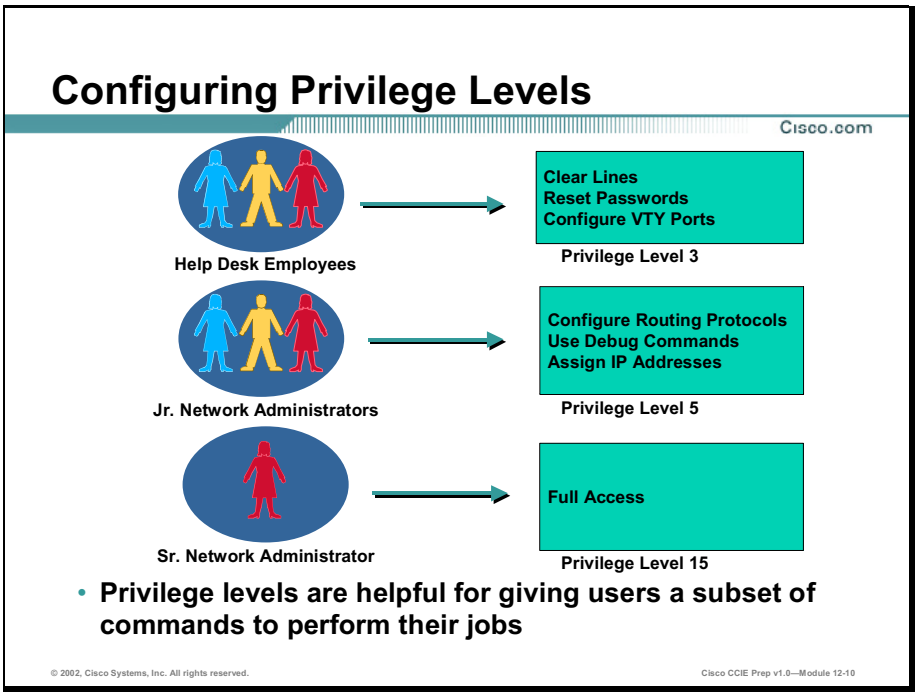

By default, the Cisco IOS has two modes: user EXEC mode and privileged (enable) mode. You can, however, configure up to 16 hierarchical levels of access in the Cisco IOS. You can then assign either usernames or passwords to access each level. By default, user mode is level 1 and privileged mode is level 15.

Privilege levels are helpful for giving non-privileged level users access to a subset of commands to perform their jobs. For example, suppose you have users who work at the help desk for an Internet Service Provider (ISP). Users may need to go in and clear lines on the network access server from time to time. You have two options here; you can give them privileged mode access to the access server or you can assign the **clear line** command to a lower privilege level and give those users access to that privilege level.

Note Level 0 is actually a non-default setting that you may want to assign to a user who only requires basic router access. By default, Level 0 can only access the following Exec commands:

<1-99>	Session number to resume
disable	Turn off privileged commands
enable	Turn on privileged commands
exit	Exit from the EXEC
help	Description of the interactive help system
logout	Exit from the EXEC

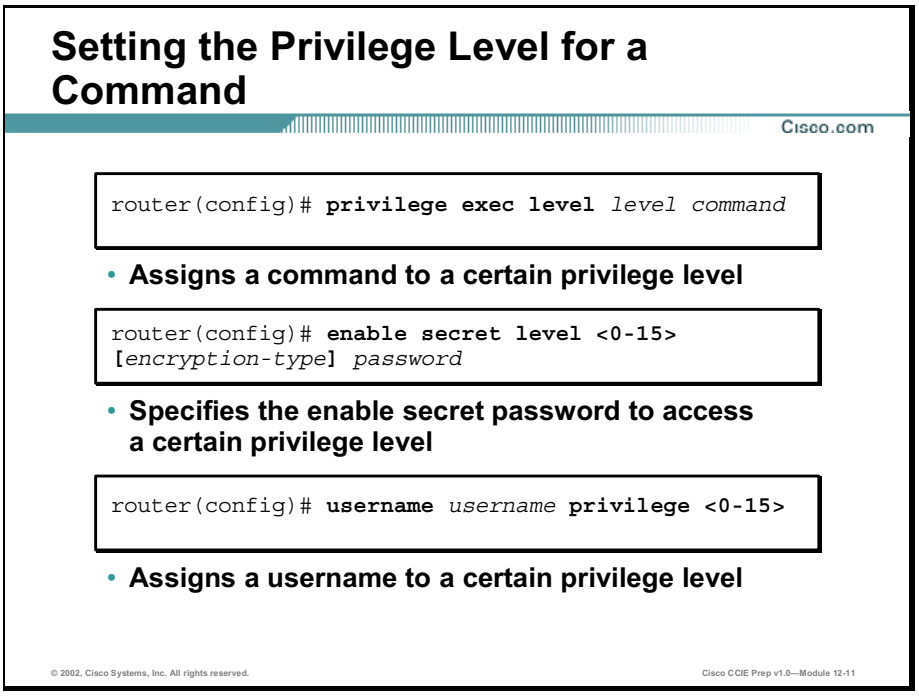

By default, all user mode commands are available at privilege levels 1 through 15 and all privileged mode commands are only available at privilege level 15. To assign a command to a different privilege level, use the following commands in global configuration mode:

Table 12-5: < privilege exec level *level command* > Command

Step	Command	Description
Step 1	**privilege exec level** *level command*	Assigns a command to a certain privilege level

In order for users to take advantage of privilege levels, you must assign an **enable password** or **enable secret** to that privilege level.

Step 2	**enable secret level** <0-15> [*encryption-type*] *password*	Specifies the enable secret password to access a certain privilege level

You can also assign users to privilege levels based on usernames in the local user database.

Step 3	**username** *username* **privilege** <0-15>	(Optional) Assigns a username to a certain privilege level

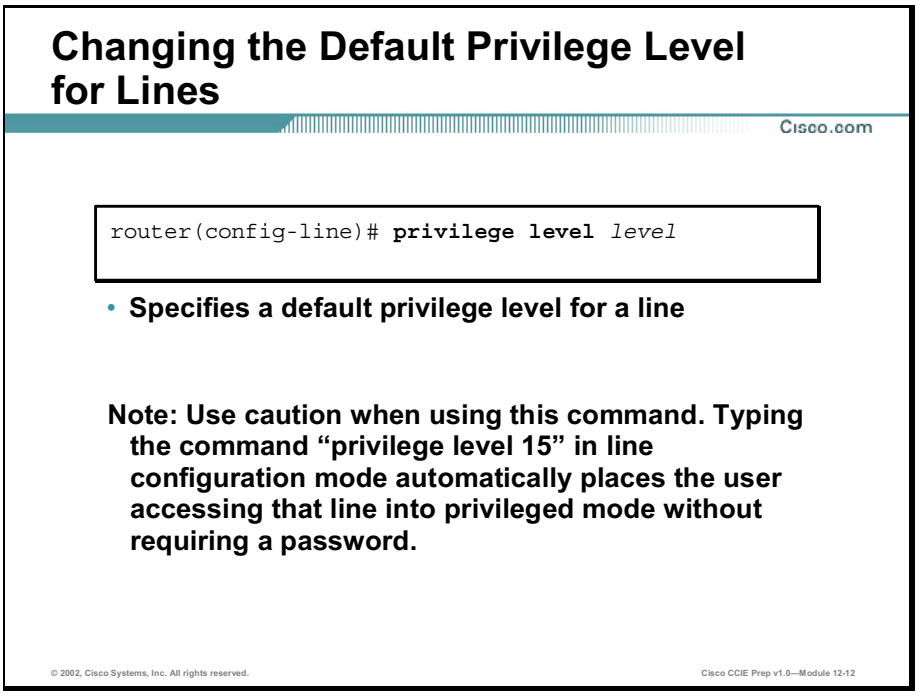

When you access the router via the console, aux, VTY, or TTY lines, you are automatically entered into user mode if a password is not assigned to the line. To change the default privilege level for a given line, use the following command in line configuration mode:

Table 12-6: <privilege level *level* > Command

Command	Description
privilege level *level*	Specifies a default privilege level for a line

A common example of this command is **privilege level 15**, which automatically places the user into privileged mode without requiring the user to enter **enable password** or **enable secret**.

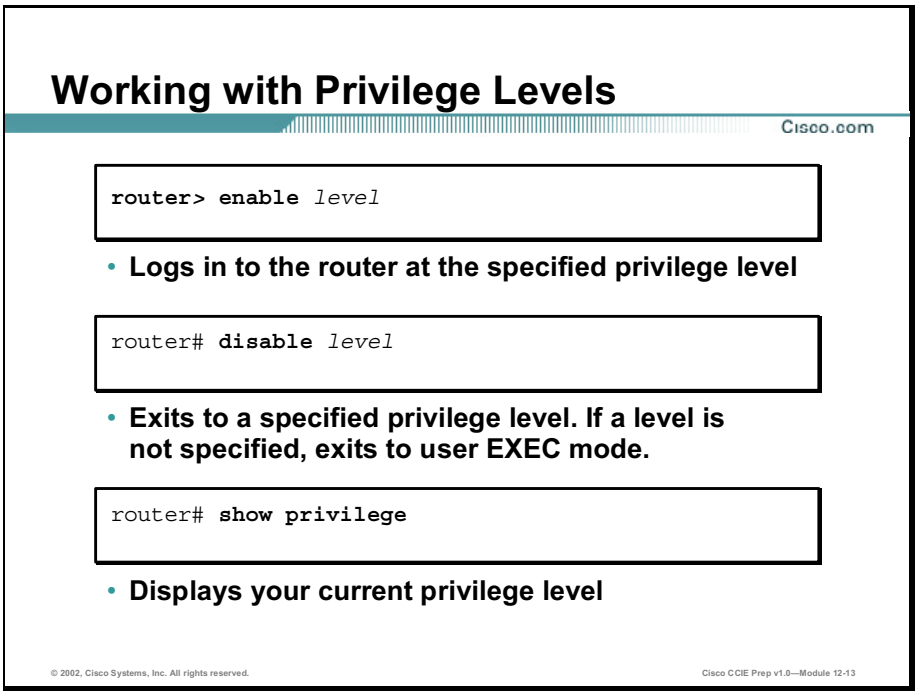

To log in to a router at a specified privilege level, use the following command in user EXEC mode:

Table 12-7: <enable *level* > Command

Command	Description
enable *level*	Logs in to the router at the specified privilege level.

To exit a specified privilege level, use the following command:

Table 12-8: <disable *level* > Command

Command	Description
disable *level*	Exits to a specified privilege level. If a level is not specified, exits to user EXEC mode.

To display your current privilege level, use the following command in EXEC mode:

Table 12-9: < show privilege > Command

Command	Description
show privilege	Displays your current privilege level.

Privilege Level Configuration Examples

This section provides privilege level configuration examples.

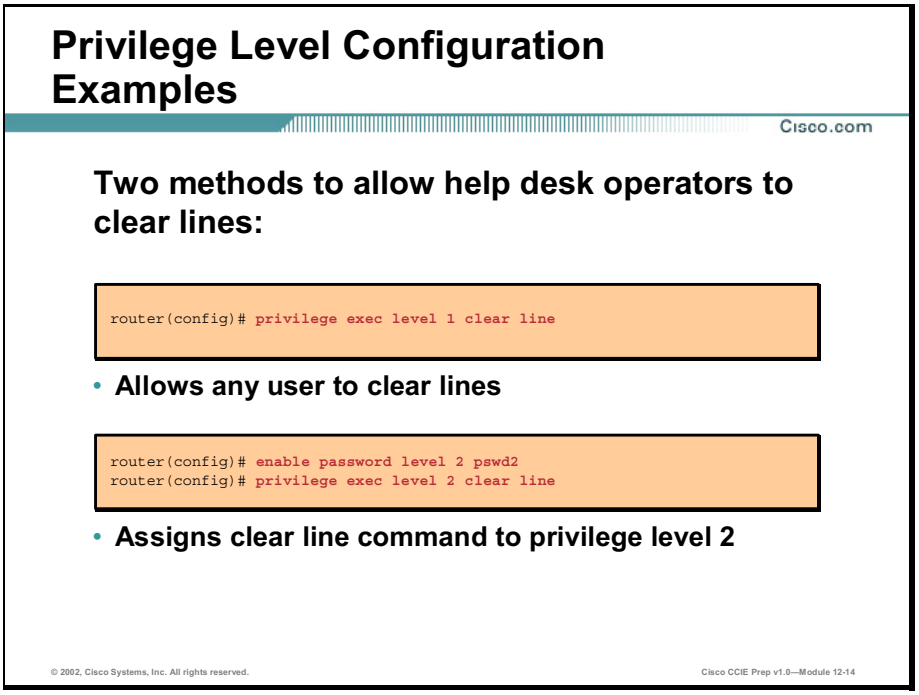

Suppose you wanted help desk operators at an ISP to be able to clear lines on the access server without giving them privileged mode access to the access server. Shown here are two examples of how this can be done using privilege levels.

If you want to allow help desk operators to clear lines, you can do either of the following:

- Change the privilege level for the **clear** and **clear line** commands to 1 (user exec mode). This allows any user to clear lines.

```
privilege exec level 1 clear line
```

- Change the privilege level for the **clear** and **clear line** commands to level 2. Then, define an **enable password** for privilege level 2, and advise only those users who need to know what the level 2 password is.

```
enable password level 2 pswd2
privilege exec level 2 clear line
```

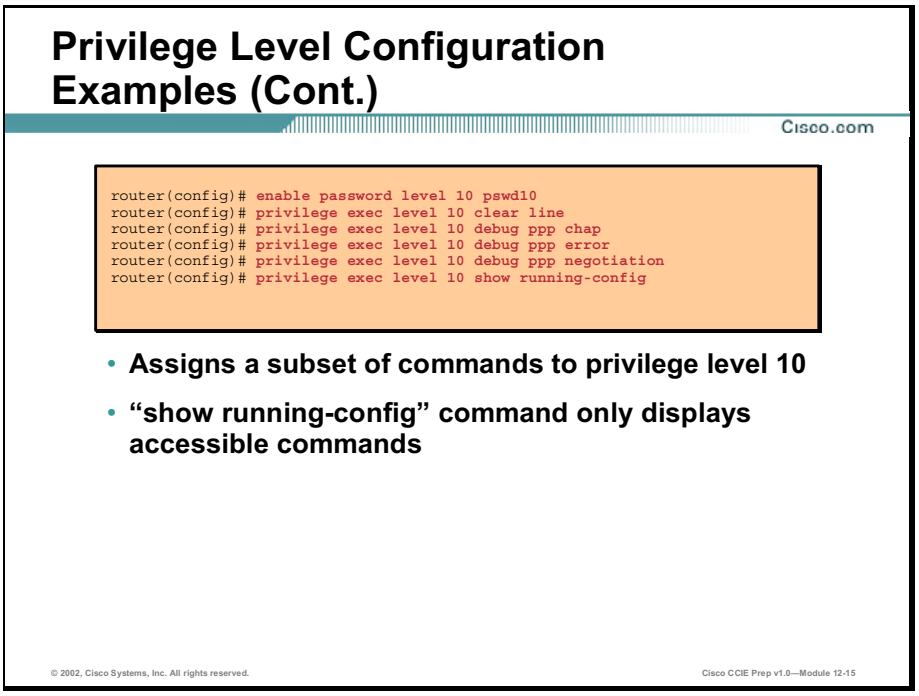

Another common example of using privilege levels is to give junior network administrators or power users access to a limited subset of privileged mode commands to perform their jobs.

In the following example, an **enable password** has been defined for privilege level 10 so that overnight system operators can use a limited number of **clear** and **debug** commands.

```
enable password level 10 pswd10
privilege exec level 10 clear line
privilege exec level 10 debug ppp chap
privilege exec level 10 debug ppp error
privilege exec level 10 debug ppp negotiation
privilege exec level 10 show running-config
```

Take note of the last command in the example. The **show running-config** command has been assigned to privilege level 10. This should allow users who access privilege level 10 with the correct **enable password** to view the running-configuration on the router. This is a common mistake. Users who access the router at privilege level 10 will be able to issue the **show running-config command**, but for security reasons, the only lines in the running-config that will be shown are lines that were created by commands that have a privilege level of 10 or lower.

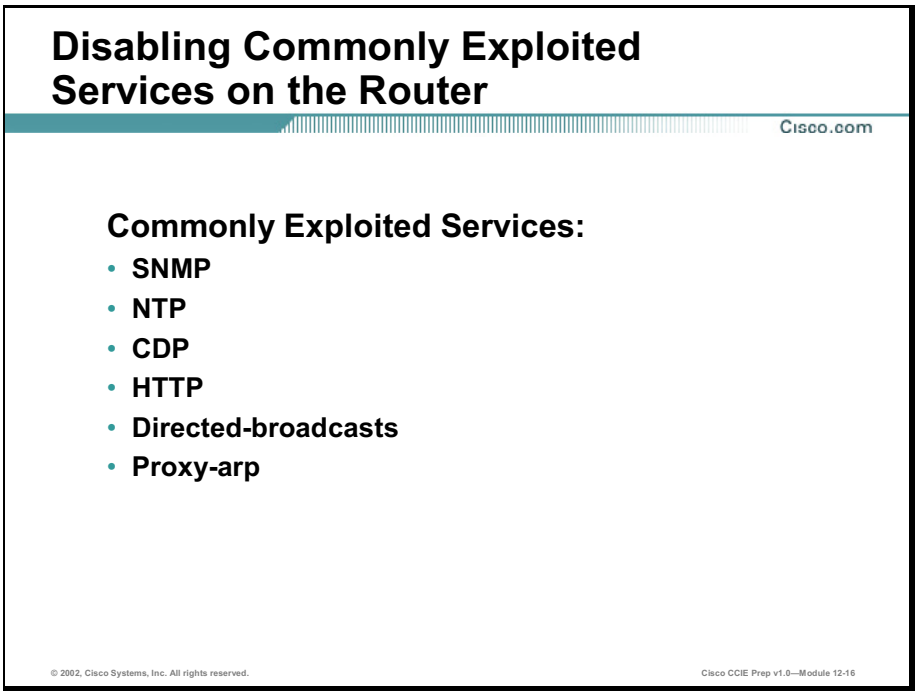

Another very important, but often overlooked, aspect of security, is disabling unused or often exploited services on the router. Here is a list of services that should be disabled if not in use on a Cisco router.

- **Simple Network Management Protocol (SNMP):** Not enabled by default. If SNMP is running on the router and you wish to disable it, use the **no snmp-server** command.
- **Network Time Protocol (NTP):** Enabled by default and can only be disabled on a per interface basis using the **ntp disable** command. If you must run NTP, configure NTP only on required interfaces, and configure NTP to listen only to certain peers using authentication or the **ntp access-group** command.
- **Cisco Discovery Protocol (CDP):** Enabled by default. Can be disabled globally with the **no cdp run** command or on a per interface basis with the **no cdp enable** command.
- **Hypertext Transport Protocol (HTTP) services**: Disabled by default. If the HTTP server is running on the router, and you wish to disable it, use the **no ip http server** command.
- You should also disable source routing. For Internet Protocol (IP), enter the **no ip source-route** global configuration command. Disabling source routing helps to prevent spoofing attacks.
- Prevent DoS smurf attacks by disabling directed broadcasts on all interfaces. For IP, use the **no ip directed-broadcast** command on each interface.
- Configure the **no ip proxy-arp** command on interfaces that connect to external networks to prevent internal addresses from being revealed. This is extremely important if you are not using Network Address Translation (NAT) to hide internal addresses from the outside world.

Access Control Lists

This section describes the operation of access lists and how they may be used to secure a network. This section also covers the differences between the various types of access lists available and describes the situations in which each type should be used.

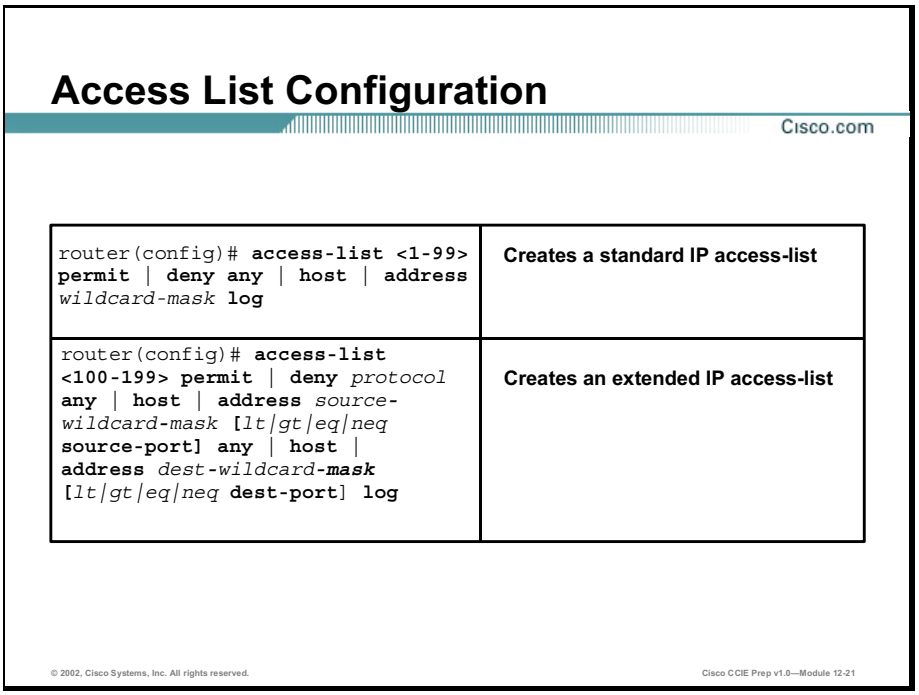

The first step is to create an access list. The second step is to apply the access list to an interface.

To create an access list, follow these steps: specify the protocol to filter, assign a unique name or number to the access list, and define packet-filtering criteria (**permit** and **deny** statements). A single access list can have multiple filtering criteria statements.

Table 12-10: Filtering Criteria Statement Commands

Command	Description
access-list <1-99> permit|**deny any**|**host**|**address** *wildcard-mask* log	Creates a standard IP access list
access-list <100-199> permit|**deny** *protocol* **any**|**host**|**address** *source-wildcard-mask* [*lt*|*gt*|*eq*|*neq* source-port] **any**|**host**|**address** *dest-wildcard-mask* [*lt*|*gt*|*eq*|*neq* dest-port] log	Creates an extended IP access list

At the end of every access list is an implicit "deny all" criteria statement. Therefore, if a packet does not match any of your criteria statements, the packet will be blocked.

Note	When creating access lists, make sure to include permit statements for routing protocol traffic. If you fail to do so, you might effectively lose communication with your neighbors, because the implicit “deny all” statement at the end of the access list will block routing updates and hello packets.

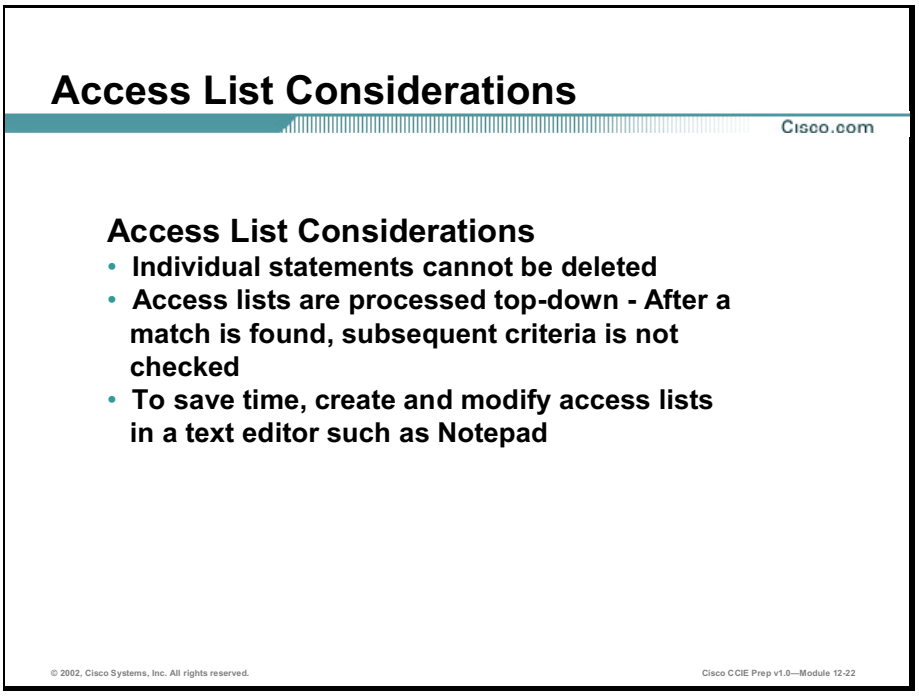

Access list entries are created in the order in which they are entered. After the access list has been applied to any interface, any additional entries that you create are appended to the end of the access list. Also note that you cannot delete individual statements after they have been created. You can only delete an entire access list.

The order that access list statements are entered is important. When the router is deciding whether to forward or block a packet, the Cisco IOS software tests the packet against each criteria statement in the order that the statements were entered. After a match is found, no more criteria statements are checked. The generally accepted rule is to put more specific statements at the top of the access list and more general statements at the bottom.

Because you cannot reorder or delete entries within an access list on a router, it is recommended that you create your access lists in a text editor, such as Notepad or TextPad, and then cut and paste this configuration into the router. This will reduce typos and save time in the CCIE lab as access lists must be removed and recreated to make changes.

If you want to make changes to an existing access list, the access list should be copied to Notepad and edited from there. You can then add the **no access-list <#>** command to the top of the file in the text editor and copy the changes to the router. This will delete the previous access list and create a new access list based on the new entries.

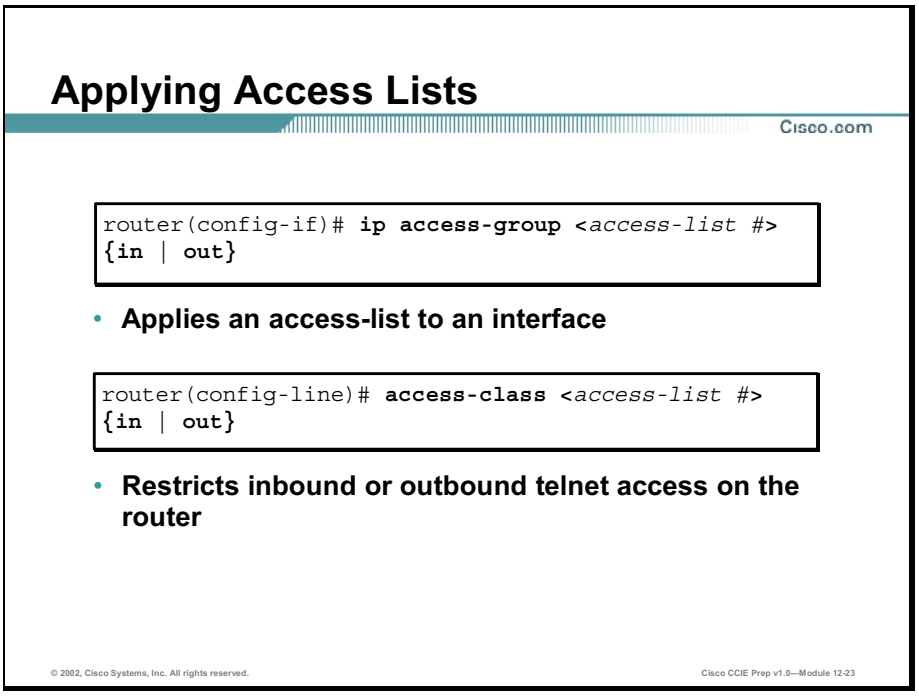

Access lists can be applied to the router in various ways. The most common examples are to an interface, to control the flow of traffic; or to VTY lines, to control Telnet access to the router.

When access lists are applied to interfaces, they can be applied in either the inbound or outbound direction. The default direction is outbound.

If the access list is applied inbound, when the router receives a packet, the Cisco IOS software checks the access list's criteria statements for a match. If the packet is permitted, the software continues to process the packet. If the packet is denied, the software discards the packet.

If the access list is applied outbound, after receiving and performing the route lookup on a packet, the packet is switched to the outbound interface. The software then checks the access list's criteria statements for a match. If the packet is permitted, the software transmits the packet. If the packet is denied, the software discards the packet.

Based on the above information, you can see that placing an access list in the outbound direction is more processor intensive than placing the access list in the inbound direction.

Use the following command in interface configuration mode to apply an access list to an interface.

Table 12-11: < ip access-group <*access-list #*> in | out > Command

Command	Description
ip access-group <*access-list #*> {**in** | **out**}	Applies an access list to an interface

Access lists can also be used to restrict Telnet access to the router. To restrict Telnet access to the router, enter the following command in VTY line configuration mode.

Table 12-12: < access-class <*access-list #*> in | out > Command

Command	Description
access-class <*access-list #*> **in** | **out**	Restricts inbound or outbound Telnet access on the router

Note Even though the focus of this section was on IP access lists, the guidelines discussed in this lesson apply to all protocols. The specific instructions for creating access lists and applying them to interfaces vary from protocol to protocol, see appendix C for specific information on access lists for each protocol.

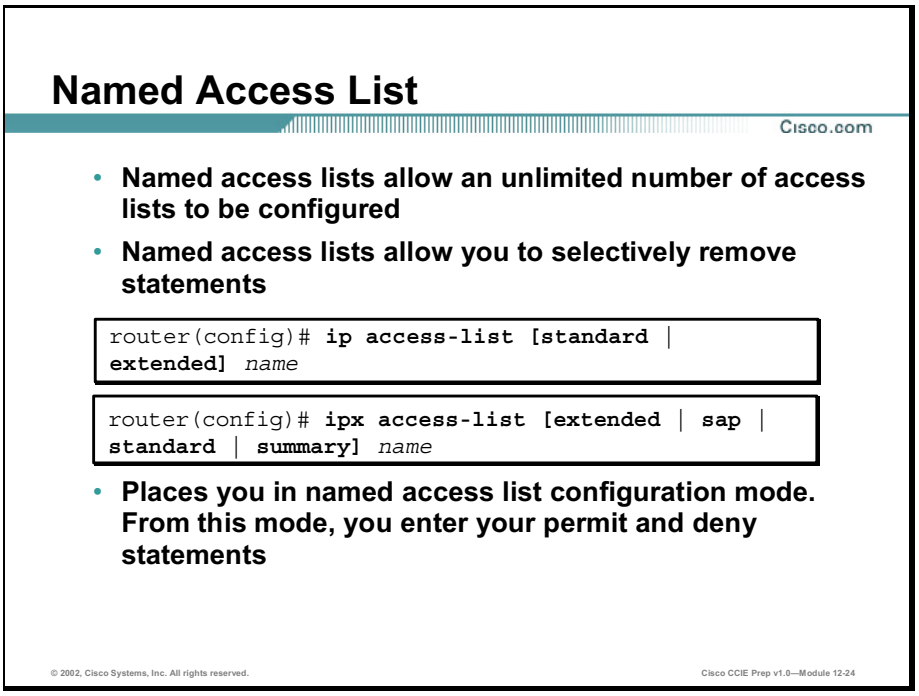

You can identify IP access lists with an alphanumeric string (a name) rather than a number. Named access lists allow you to configure an unlimited number of access lists, whereas numbered access lists are limited to certain number ranges. Another advantage to named access lists is that they allow you to selectively remove entries from the access list. As covered in the previous section, numbered access lists require you to delete the access list and rebuild it to delete any entries.

Table 12-13: IP Access-List Commands

Command	Description
ip access-list [**standard** | **extended**] *name*	This command places you in named access list configuration mode. From this mode you enter your **permit** and **deny** statements.

Consider the following guidelines before configuring named access lists:

- Access lists specified by name are not compatible with Cisco IOS Releases prior to 11.2.
- Named access lists are available for standard and extended IP access lists.
- A standard access list and an extended access list cannot have the same name.

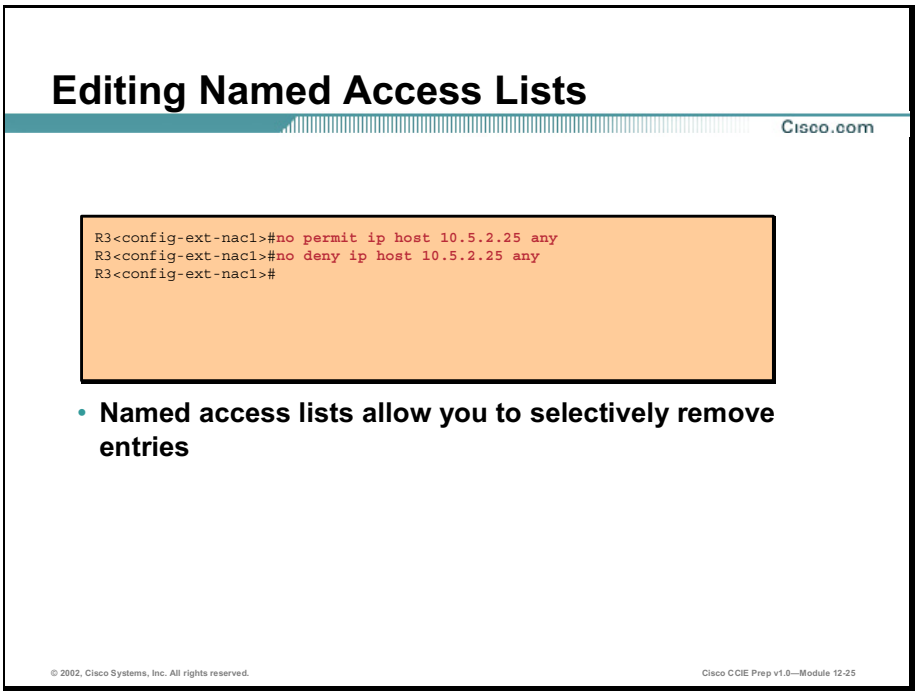

After you create an access list (named or numbered), you can place subsequent permit and deny entries at the end of the list only. In other words, you cannot selectively add permit or deny entries to an existing access list. However, named access lists offer the advantage of allowing you to selectively remove entries using the **no permit** and **no deny** commands.

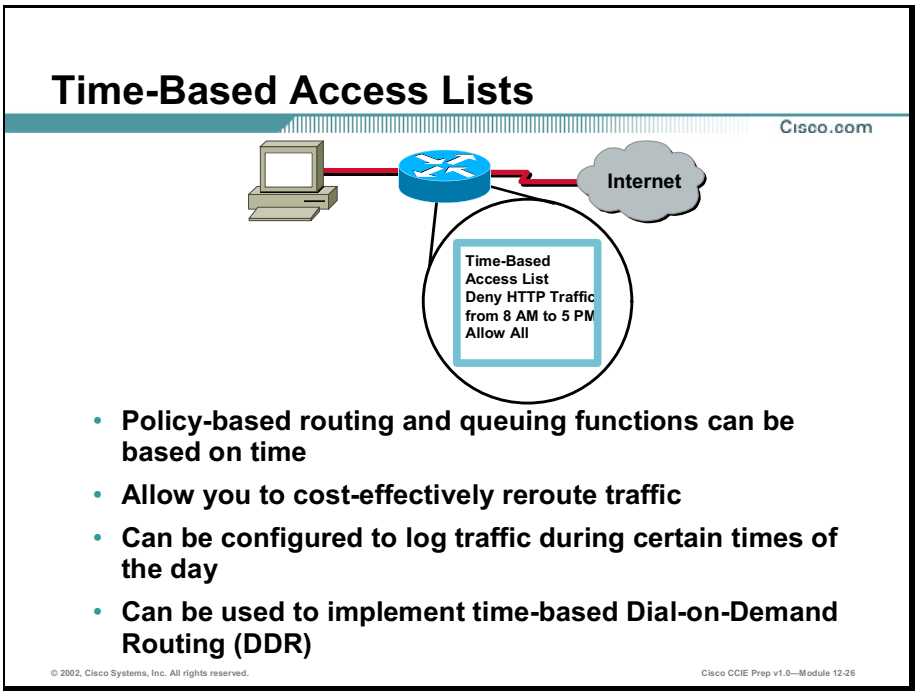

One of the problems associated with using normal access lists is that they are in effect from the moment they are applied to an interface. Time-based access lists are used to restrict or allow access to certain resources based on a time range. Time ranges can be based on time of the day, day of the week, or an absolute period of time, based on a start and end time. There are many possible benefits of using time-based access lists. Some of these benefits are listed below.

- Allows network administrators to set time-based security policies, including:
  - Perimeter security using Context-Based Access Control (CBAC) or access lists
  - Data confidentiality (time-based crypto access lists) with IP Security (IPSec)
- Policy-based routing and queuing functions can be based on time.
- If you have multiple service providers and their rates vary by time of day, time-based access lists allow you to cost-effectively reroute traffic.
- Allows service providers to dynamically change Committed Access Rate (CAR) configurations to support the Quality of Service (QoS) Service Level Agreements (SLAs) that are negotiated for different periods throughout the day.
- Access list entries can be configured to log traffic during certain times of the day, allowing network administrators to control the amount of logging messages received.
- Can be used to implement time-based Dial-on-Demand Routing (DDR).

Note Time ranges rely on the router's system clock. For this feature to work reliably, it is recommended that you use Network Time Protocol (NTP) to synchronize the router's clock with an Internet time source.

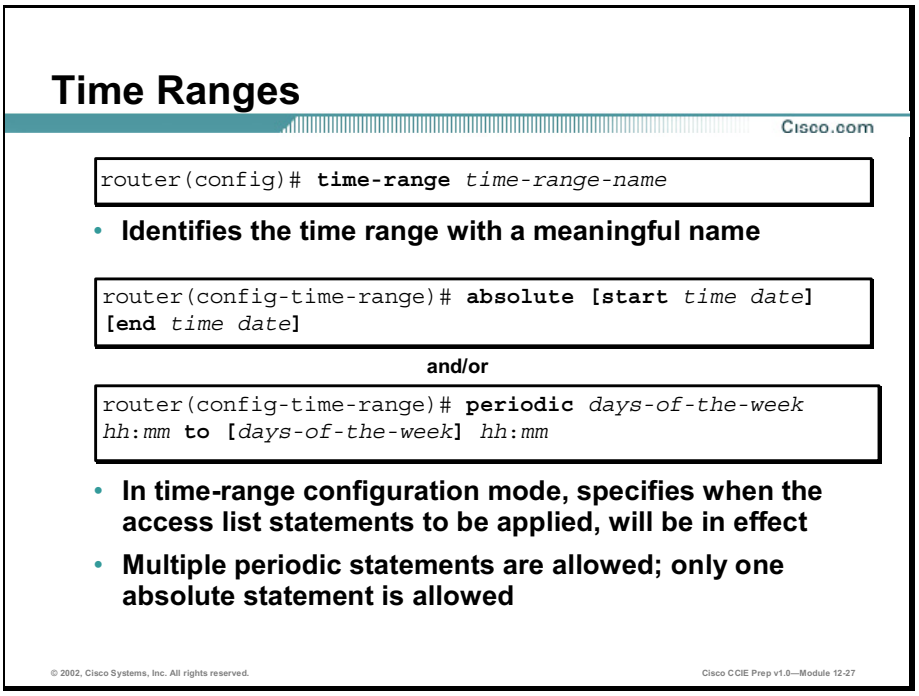

Currently, extended access lists are the only functions that can use time ranges. The time range allows the network administrator to define when the **permit** or **deny** statements in the access list are in effect. Time ranges can be based on a recurring **periodic** time period or an **absolute** start and end time.

To define a time range, use the following commands beginning in global configuration mode.

Table 12-14: Define Time Range Commands

Step	Command	Description
Step 1	**time-range** *time-range-name*	Identifies the time range with a meaningful name
Step 2	**absolute** [**start** *time date*] [**end** *time date*] and/or **periodic** *days-of-the-week hh:mm* **to** *[days-of-the-week] hh:mm*	In time-range configuration mode, specifies when the access list statements to be applied, will be in effect. Multiple **periodic** statements are allowed; only one **absolute** statement is allowed. **absolute start** *time date* - Absolute time and date that the associated **permit** or **deny** statement goes into effect. The *time* is expressed in a 24-hour clock, in the form of *hours:minutes*. For example, 8:00 is 8:00 am and 20:00 is 8:00 pm. The *date* is expressed in the format *day month year*. The minimum start is 00:00 1 January 1993. If no start time and date are specified, the **permit** or **deny** statement is in effect immediately. **end** *time date* - Absolute time and date that the associated **permit** or **deny** statement is no longer in effect. Same *time* and *date* format as described for the **start**. The end time and date must be after the start time and date. The maximum end time is 23:59 31 December 2035. If no end time and date are specified, the **permit** or **deny** statement is in effect indefinitely. **periodic** *days-of-the-week* - The first occurrence of this argument is the starting day or days that the associated time range is in effect. The second occurrence is the ending day or days the associated statement is no longer in effect. This argument can be any single day or combinations of days: **Monday**, **Tuesday**, **Wednesday**, **Thursday**, **Friday**, **Saturday**, and **Sunday**. Other possible values are: **daily** -- Monday through Sunday **weekdays** -- Monday through Friday **weekend** -- Saturday and Sunday *hh:mm* - The first occurrence of this argument is the starting *hours***:***minutes* when the associated time range is in effect. The second occurrence is the ending *hours***:***minutes* when the associated statement is no longer in effect.

Repeat these tasks if you have multiple items that you want in effect at different times. For example, repeat the steps to include multiple **permit** or **deny** statements in an access list to be in effect at different times.

Note If a **time-range** command has both **absolute** and **periodic** values specified, then the **periodic** items are evaluated only after the **absolute start** time is reached, and they are no longer evaluated after the **absolute end** time is reached.

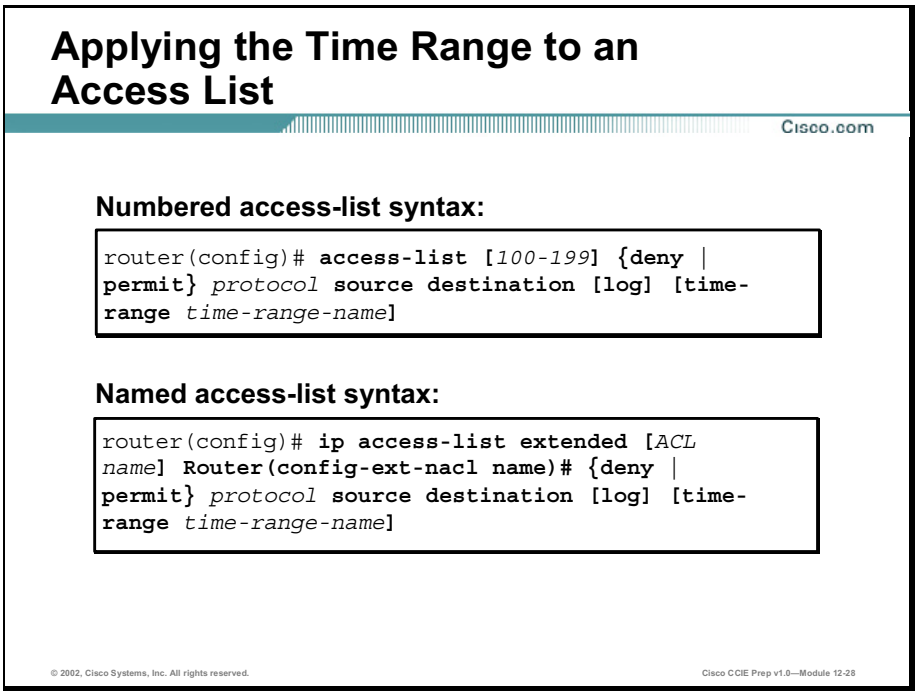

Numbered access list example:

```
router(config)# access-list [100-199] {deny | permit} protocol source
destination [log] [time-range time-range-name]
```

Named access list example:

```
router(config)# ip access-list extended [ACL name]
router(config-ext-nacl)# {deny | permit} protocol source destination [log]
[time-range time-range-name]
```

Once the time range has been configured, it needs to be referenced in an access list. Both named and numbered access lists can reference a time range. Permit and deny statements in an access list referencing a time range will only apply during that time range.

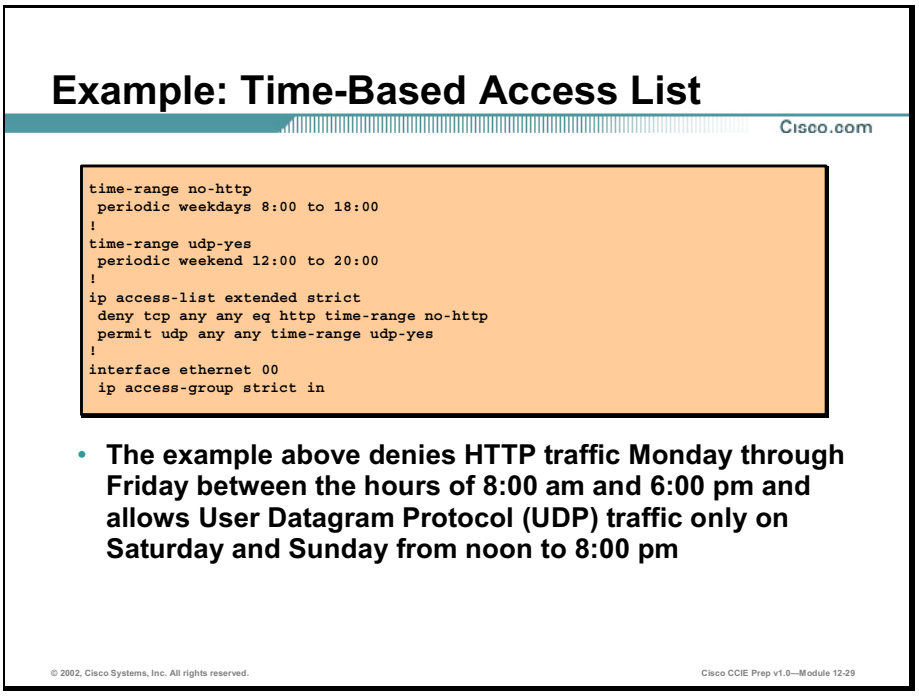


```
time-range no-http
 periodic weekdays 8:00 to 18:00
!
time-range udp-yes
 periodic weekend 12:00 to 20:00
!
ip access-list extended strict
 deny tcp any any eq http time-range no-http
 permit udp any any time-range udp-yes
!
interface ethernet 0/0
 ip access-group strict in
```

The example above denies HTTP traffic Monday through Friday between the hours of 8:00 am and 6:00 pm and allows User Datagram Protocol (UDP) traffic only on Saturday and Sunday from noon to 8:00 pm.

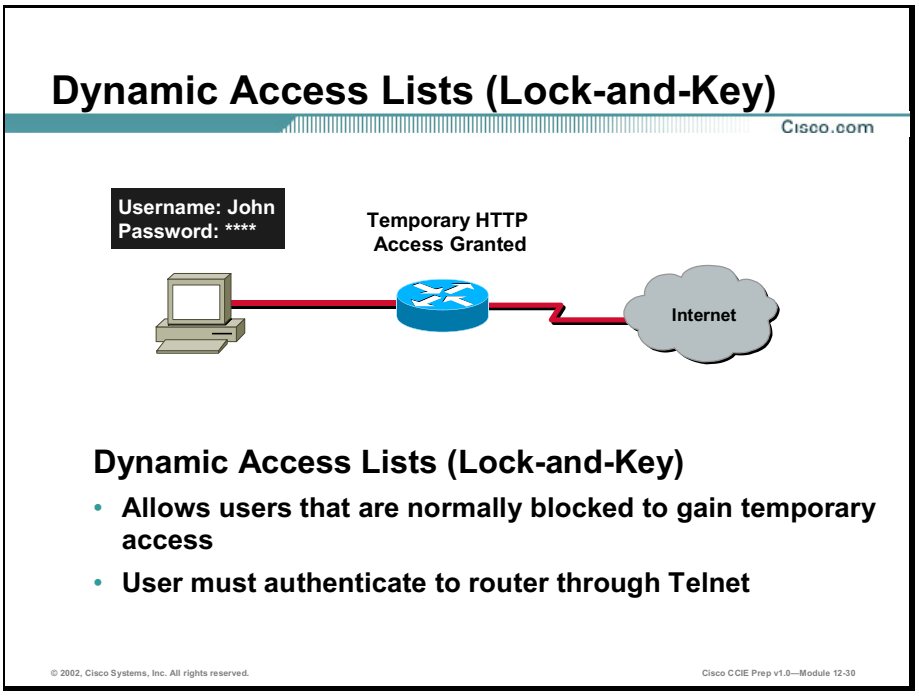

Lock-and-Key is a traffic filtering security feature that dynamically filters IP protocol traffic. Lock-and-Key is configured using dynamic extended IP access lists. It can be used in conjunction with other standard access lists and static extended access lists.

When Lock-and-Key is configured, designated users whose IP traffic is normally blocked at a router by an inbound access list can gain temporary access through the router. When triggered, Lock-and-Key reconfigures the interface's existing IP access list to permit authenticated outside users to reach designated inside resources. After a period of time, Lock-and-Key reconfigures the access list back to its original state.

For a user to gain access to an inside resource through a router configured for Lock-and-Key, the user must first Telnet to the router. When a user initiates a standard Telnet session to the router, Lock-and-Key automatically attempts to authenticate the user. If the user is authenticated successfully, the user will then gain temporary access through the router to the inside network.

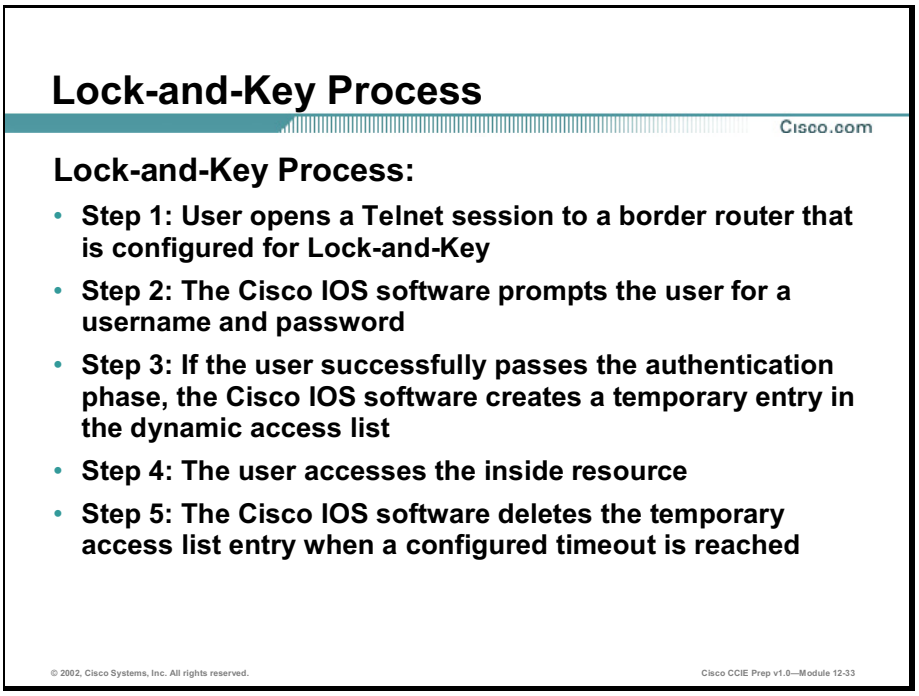

1. User opens a Telnet session to a border router that is configured for Lock-and-Key.
2. The Cisco IOS software receives the Telnet packet, allows the Telnet session, and prompts the user for a username and password.
3. If the user successfully passes the authentication phase, they are automatically logged out of the Telnet session, and the Cisco IOS software creates a temporary entry in the dynamic access list.
4. The user accesses the inside resource.
5. The Cisco IOS software deletes the temporary access list entry when a configured timeout is reached, or when the system administrator manually clears the dynamic entry. The configured timeout can either be an idle timeout or an absolute timeout.

Note	The temporary access list entry is not automatically deleted when the user terminates a session. The temporary access list entry remains until a configured timeout is reached or until it is cleared by the system administrator.

To manually delete a temporary access list entry, perform the following task in privileged EXEC mode:

Table 12-15: < clear access-template [*access-list-number* | *name*] [*dynamic-name*] [*source*] [*destination*] **> Command**

Command	Description
clear access-template [*access-list-number* | *name*] [*dynamic-name*] [*source*] [*destination*]	Manually deletes a dynamic access list entry

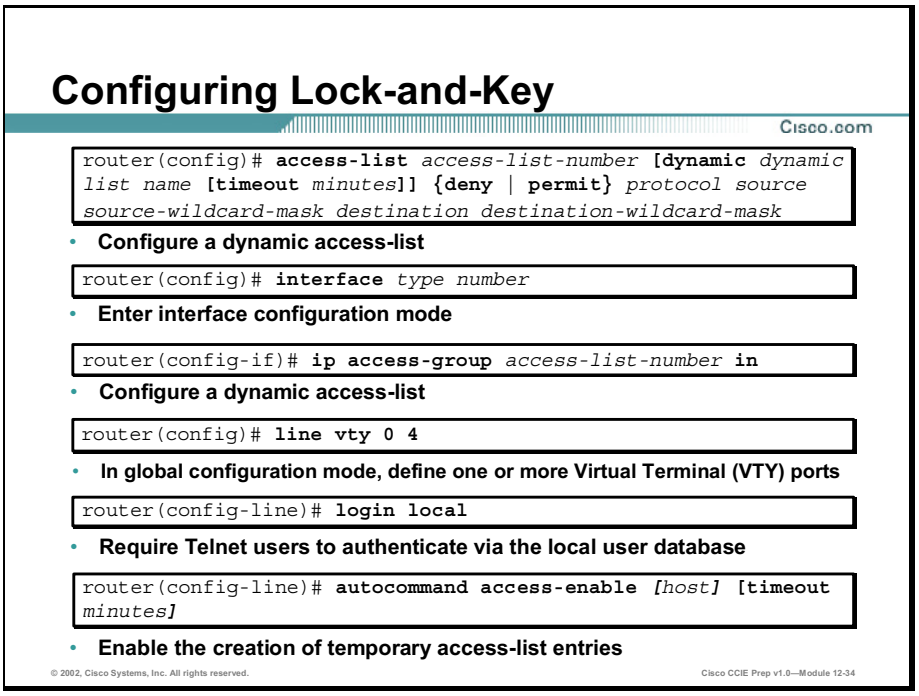

Here are commands required in order to configure Lock-and-Key:

Table 12-16: Lock-and-Key Commands

Step	Task	Command
Step 1	Configure a dynamic access list. The dynamic access list serves as a template and placeholder for temporary access list entries. Remember that the source and source-wildcard are always replaced with the IP address of the authenticating host, so it is a good security practice to use the keyword **any** for the source IP address of your dynamic entry.	**access-list** access-list-number [**dynamic** *dynamic list name* [**timeout** minutes]] {**deny** | **permit**} **protocol** *source source-wildcard-mask destination destination-wildcard-mask*
Step 2	Enter interface configuration mode.	**interface** *type number*
Step 3	Apply the access list to the interface.	**ip access-group** *access-list-number* **in**
Step 4	In global configuration mode, define one or more Virtual Terminal (VTY) ports. If you specify multiple VTY ports, they must all be configured identically because the software hunts for available VTY ports on a round-robin basis. If you do not want to configure all your VTY ports for Lock-and-Key access, you can specify a group of VTY ports for Lock-and-Key support only.	**line vty 0 4**
Step 5	Require Telnet users to authenticate via the local user database	**login local**
Step 6	Enable the creation of temporary access list entries. If the host argument is not specified, all hosts on the entire network are allowed to set up a temporary access list entry. The dynamic access list contains the network mask to enable the new network connection.	**autocommand access-enable** [**host**] [**timeout** *minutes*]

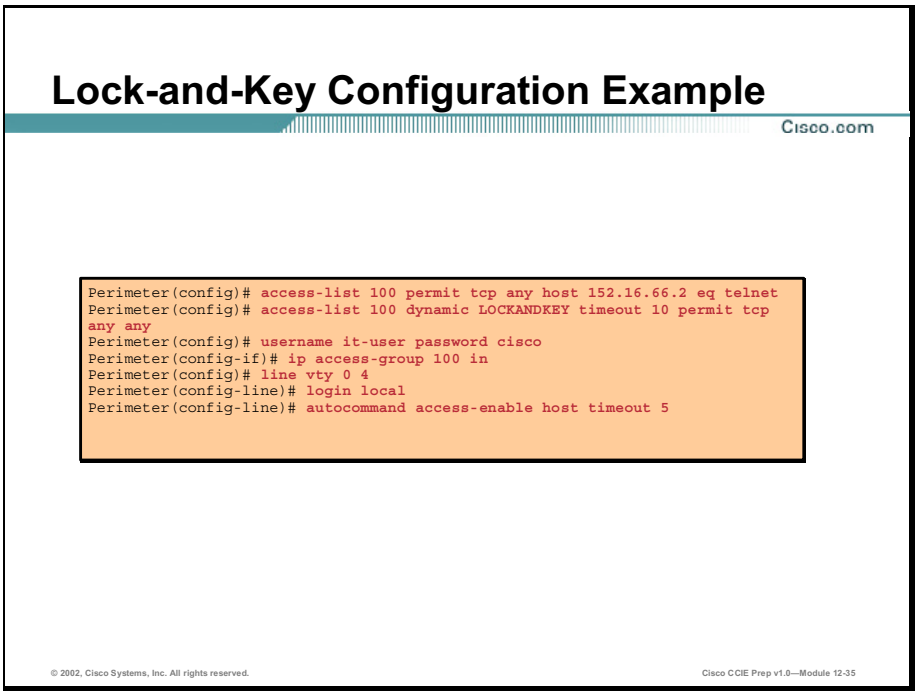


```
Perimeter(config)# access-list 100 permit tcp any host 152.16.66.2 eq telnet
Perimeter(config)# access-list 100 dynamic LOCKANDKEY timeout 10 permit tcp any any
Perimeter(config)# username it-user password cisco
Perimeter(config-if)# ip access-group 100 in
Perimeter#(config)# line vty 0 4
Perimeter#(config-line)# login local
Perimeter#(config-line)# autocommand access-enable host timeout 5
```

The first line in the example above permits Telnet access to the perimeter router's interface that connects to the Internet (152.16.66.2). This line is required, otherwise you will never be able to Telnet to the perimeter router, authenticate, and create the dynamic entry.

The second line is your dynamic entry that will be inserted into access list 100 when the user is successfully authenticated. The timeout entry is called the absolute-timeout and is the amount of time you want to leave this entry open, in this case 10 minutes. You are using the local user database for authentication.

You must apply access list 100 to the Perimeter router's interface that is connected to the Internet. Currently, only the first entry in this access list is active. The first entry blocks everything and permits Telnet access only to the perimeter router.

Finally, the last step needed in order to allow creation of the dynamic entry is to use the **autocommand** command on one or more VTY lines. The **host** keyword is very important; without it the dynamic entry would not substitute the user's source address in the dynamic entry. The command **timeout 5** is optional and sets the idle time-out.

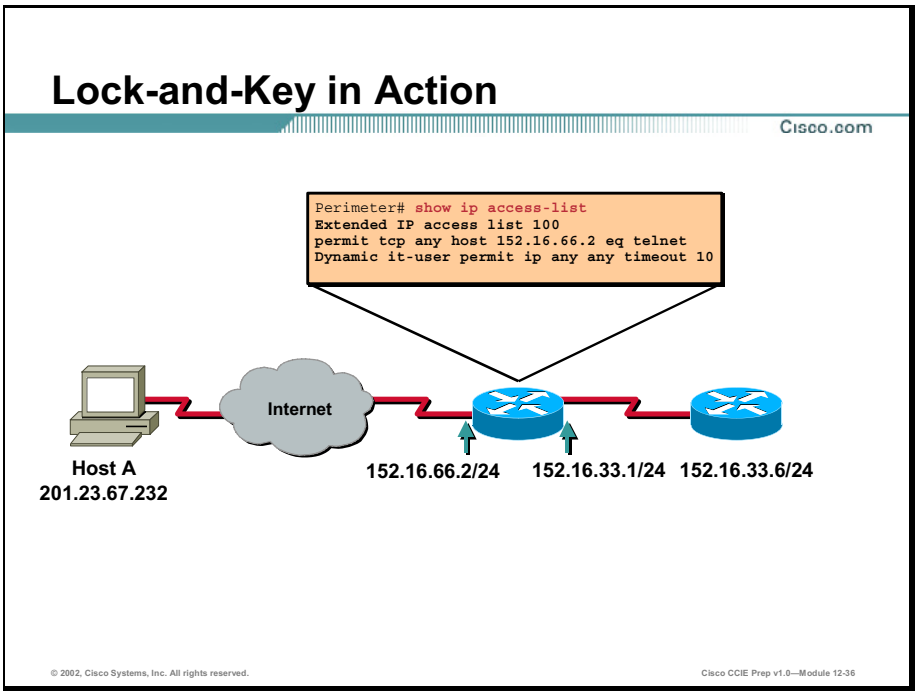

This example shows a dynamic access list in action. Our perimeter router's IP address is 152.16.66.2. The inside resource that you need to access is another router at 152.16.33.6.

Our ISP has assigned a dynamic IP address of 201.23.67.232. This is where you will be Telneting from.

First, verify the current access list entries on the Perimeter router.

```
Perimeter# show ip access-list
Extended IP access list 100
permit tcp any host 152.16.66.2 eq telnet
Dynamic it-user permit ip any any timeout 10
```

Notice that there are no entries listed under the Dynamic entry.

Now, try to Telnet to 152.16.33.6, which is the router you want to work on:

```
C:\telnet 152.16.33.6
Trying 152.16.33.6 ...
% Destination unreachable; gateway or host down
```

As expected, you are not able to Telnet directly to the inside resource.

Next, authenticate to the Perimeter router, which should create the dynamic entry.

```
telnet 152.16.66.2
Trying 152.16.66.2 ... Open
```

User Access Verification

```
Username: it-user
```

```
Password:
[Connection to 152.16.66.2 closed by foreign host]
```

After successful authentication, notice how your session was dropped. Verify that the dynamic entry was created in the access list 100.

```
Perimeter# show ip access-list 100
Extended IP access list 100
permit tcp any host 152.16.66.2 eq telnet (56 matches)
Dynamic test permit ip any any timeout 10
permit ip host 201.23.67.232 any timeout 10 (time left 439)
```

Notice that a dynamic entry is now active and your IP address of 201.23.67.232 was inserted into that entry as the source address

You should now be able to Telnet directly to the internal router 152.16.33.6.

```
telnet 152.16.33.6
Trying 152.16.33.6 ... Open
User Access Verification
Password:
R3>
```

You now have access to the internal router and can proceed with any work that needs to be done. Your connection will only stay open for the configured 10 minutes, at which time it will be automatically closed.

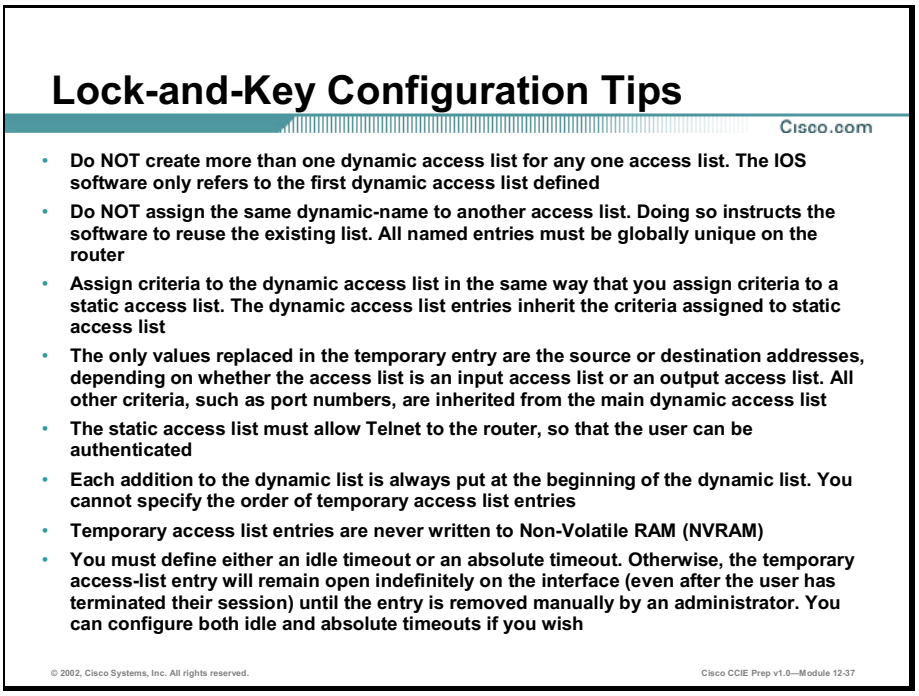

- Do NOT create more than one dynamic access list for any one access list. The IOS software only refers to the first dynamic access list defined.
- Do NOT assign the same dynamic name to another access list. Doing so instructs the software to reuse the existing list. All named entries must be globally unique on the router.
- Assign criteria to the dynamic access list in the same way that you assign criteria to a static access list. The dynamic access list entries inherit the criteria assigned to static access list.
- The only values replaced in the temporary entry are the source or destination addresses, depending on whether the access list is an input or output access list. All other criteria, such as port numbers, are inherited from the main dynamic access list.
- The static access list must allow Telnet to the router, so that the user can be authenticated.
- Each addition to the dynamic list is always put at the beginning of the dynamic list. You cannot specify the order of temporary access list entries.
- Temporary access list entries are never written to Non-Volatile RAM (NVRAM).
- You must define either an idle timeout or an absolute timeout. Otherwise, the temporary access list entry will remain open indefinitely on the interface (even after the user has terminated their session) until an administrator removes the entry manually. You can configure both idle and absolute timeouts if you wish.
- If you configure both idle and absolute timeouts, the idle timeout value must be less than the absolute timeout value.

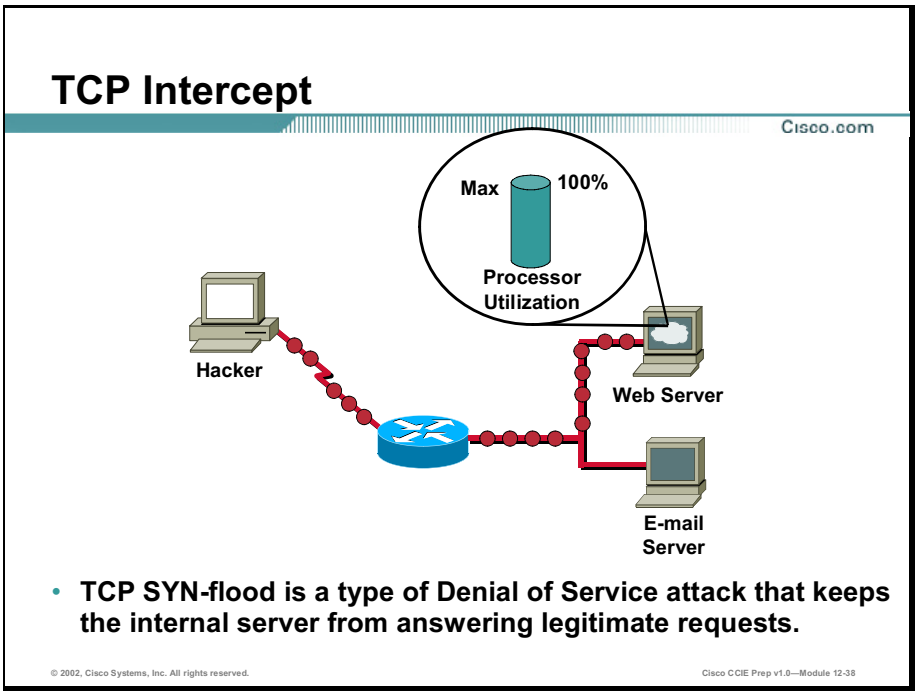

The Transmission Control Protocol (TCP) intercept feature protects internal servers from TCP Synchronization (SYN)-flood attacks, which are a type of denial-of-service attack.

A SYN-flood attack occurs when a hacker floods a server with a barrage of requests for a TCP connection. These requests are being sent from a spoofed IP address and therefore the server can never build the three-way handshake to establish the TCP connection. The resulting volume of unresolved open connections eventually overwhelms the server and can cause it to deny service to valid requests, preventing legitimate users from connecting to a web site, accessing e-mail, etc.

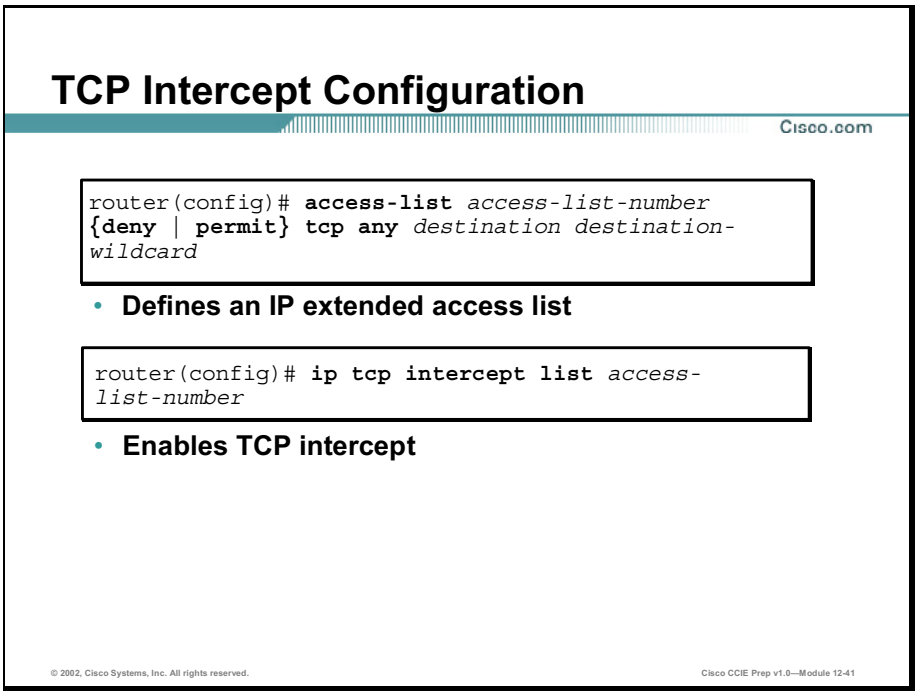


```
ip tcp intercept list 101
access-list 101 permit tcp any 192.168.1.0 0.0.0.255
```

To configure TCP intercept, perform the following tasks. The first task is required. The other tasks are optional.

- Enabling TCP Intercept (Required)
- Setting the TCP Intercept Mode (Optional)
- Setting the TCP Intercept Drop Mode (Optional)
- Changing the TCP Intercept Timers (Optional)
- Changing the TCP Intercept Aggressive Thresholds (Optional)
- Monitoring and Maintaining TCP Intercept (Optional)

Enabling TCP Intercept

You can define an access list to intercept all requests or only those coming from specific networks or destined for specific servers. Typically, the access list will define the source as **any** and define specific destination networks or servers. That is, you do not attempt to filter on the source addresses since you do not necessarily know who to intercept packets from. You identify the destination in order to protect destination servers. If no access list match is found, the router allows the request to pass with no further action.

To enable TCP intercept, use the following commands in global configuration mode:

Table 12-17: TCP Intercept Commands

Step	Command	Description
Step 1	**access-list** *access-list-number* {**deny** | **permit**} **tcp** any *destination destination-wildcard*	Defines an IP extended access list
Step 2	**ip tcp intercept list** *access-list-number*	Enables TCP intercept

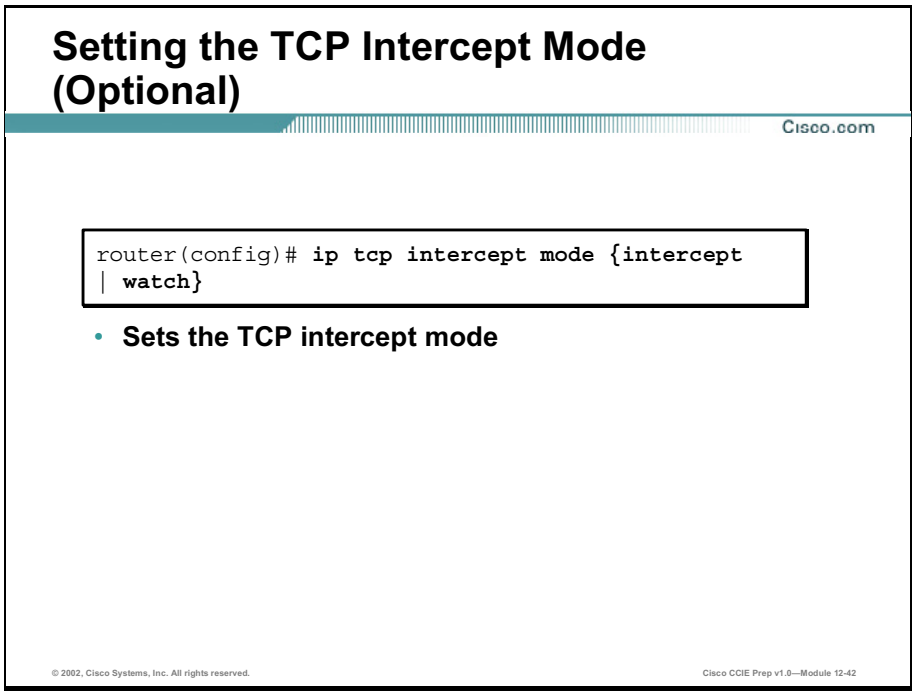

The TCP intercept can operate in either active intercept mode or passive watch mode. The default is intercept mode.

In intercept mode, the IOS software actively intercepts each incoming connection request (SYN) and responds on behalf of the server with a SYN-ACK, then waits for an Acknowledge (ACK) from the client. When that ACK is received, the original SYN is sent to the server and the IOS software performs a three-way handshake with the server. When this is complete, the two half-connections are joined.

In watch mode, connection requests are allowed to pass through the router to the server but are tracked until they become established. If they fail to become established within 30 seconds (configurable with the **ip tcp intercept watch-timeout** command), the software sends a TCP RESET to the server to clear the half-formed connection.

To set the TCP intercept mode, use the following command in global configuration mode:

Table 12-18: < ip tcp intercept mode {**intercept** | **watch**} **> Command**

Command	Description
ip tcp intercept mode {**intercept** | **watch**}	Sets the TCP intercept mode

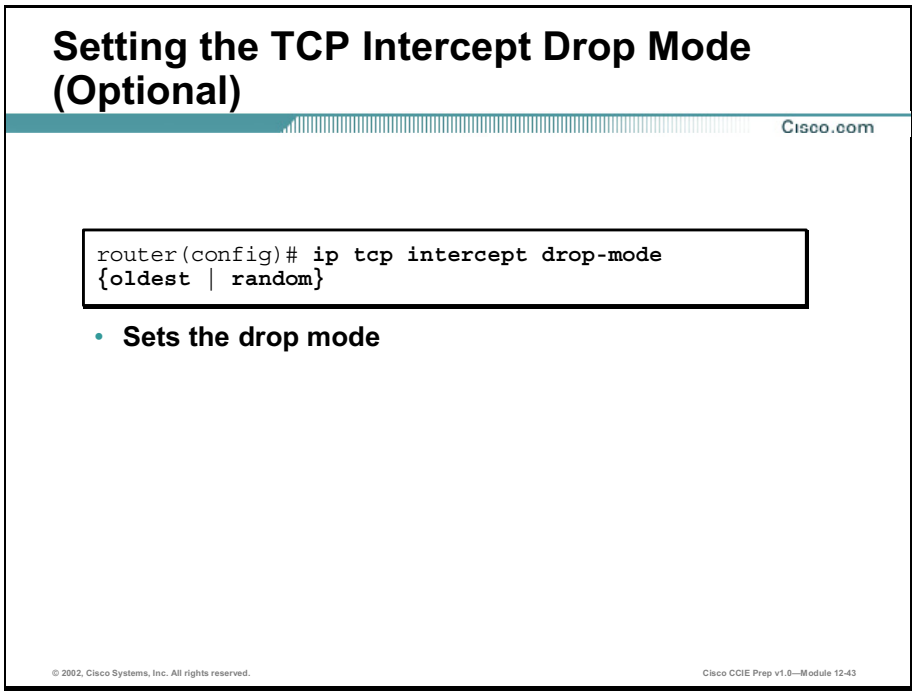

When under attack, the TCP intercept feature becomes more aggressive in its protective behavior. If the number of incomplete connections exceeds 1100 or the number of connections arriving in the last one minute exceeds 1100, each new arriving connection causes the oldest half-formed connection to be deleted. Also, the initial retransmission timeout is reduced to 0.5 seconds, so that the total time to attempt to establish a connection is cut in half.

By default, the IOS software drops the oldest half-formed connection. Alternatively, you can configure the IOS software to randomly drop connections. To set the drop mode, use the following command in global configuration mode:

Table 12-19: < ip tcp intercept drop-mode {oldest | random} > Command

Command	Description
ip tcp intercept drop-mode {**oldest** | **random**}	Sets the drop mode

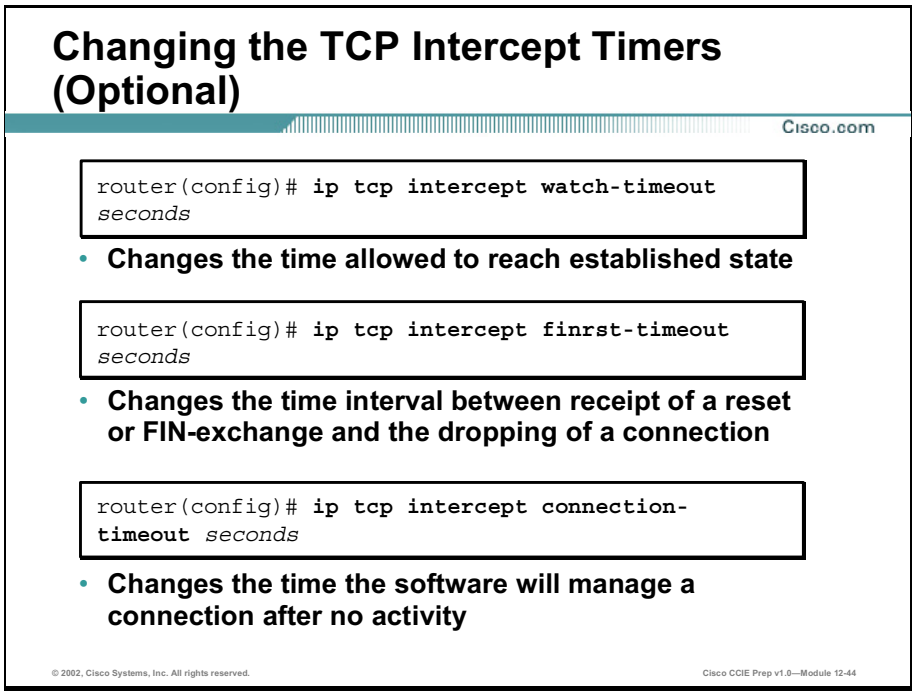

By default, the IOS software waits for 30 seconds when in watch mode for a watched connection to reach established state before sending a TCP RESET to the server. To change this value, use the following command in global configuration mode.

Table 12-20: < ip tcp intercept watch-timeout *seconds* > Command

Command	Description
ip tcp intercept watch-timeout *seconds*	Changes the time allowed to reach established state

By default, the IOS software waits for 5 seconds from receipt of a TCP RESET or FIN packet before it ceases to manage the connection. To change this value, use the following command in global configuration mode.

Table 12-21: < ip tcp intercept finrst-timeout *seconds* > Command

Command	Description
ip tcp intercept finrst-timeout *seconds*	Changes the time interval between receipt of a reset or FIN-exchange and the dropping of a connection

By default, the software still manages a connection for 24 hours after no activity. To change this value, use the following command in global configuration mode:

Table 12-22: < ip tcp intercept connection- timeout *seconds* > Command

Command	Description
ip tcp intercept connection-timeout *seconds*	Changes the time the software will manage a connection after no activity

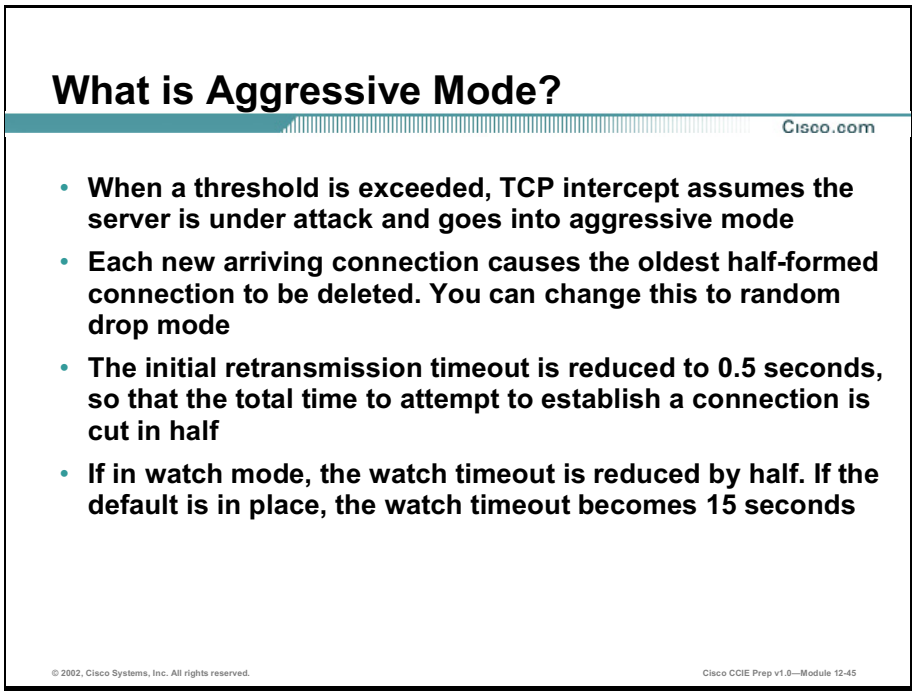

What is aggressive mode?

When a threshold is exceeded, TCP intercept assumes the server is under attack and goes into aggressive mode. When in aggressive mode, the following occurs:

- Each new arriving connection causes the oldest half-formed connection to be deleted. You can change this to random drop mode.
- The initial retransmission timeout is reduced to 0.5 seconds, so that the total time to attempt to establish a connection is cut in half. When not in aggressive mode, the IOS software exponentially backs off on its retransmissions of SYN segments. The initial retransmission timeout is 1 second. The subsequent timeouts are 2 seconds, 4 seconds, 8 seconds, and 16 seconds. The IOS software retransmits 4 times before giving up; therefore, the half-formed connection is deleted after 31 seconds of no acknowledgments received.
- If in watch mode, the watch timeout is reduced by half. If the default is in place, the watch timeout becomes 15 seconds.

Note	The two factors that determine aggressive behavior are related and work together. When either of the high values is exceeded, aggressive behavior begins. When both quantities fall below the low value, aggressive behavior ends.

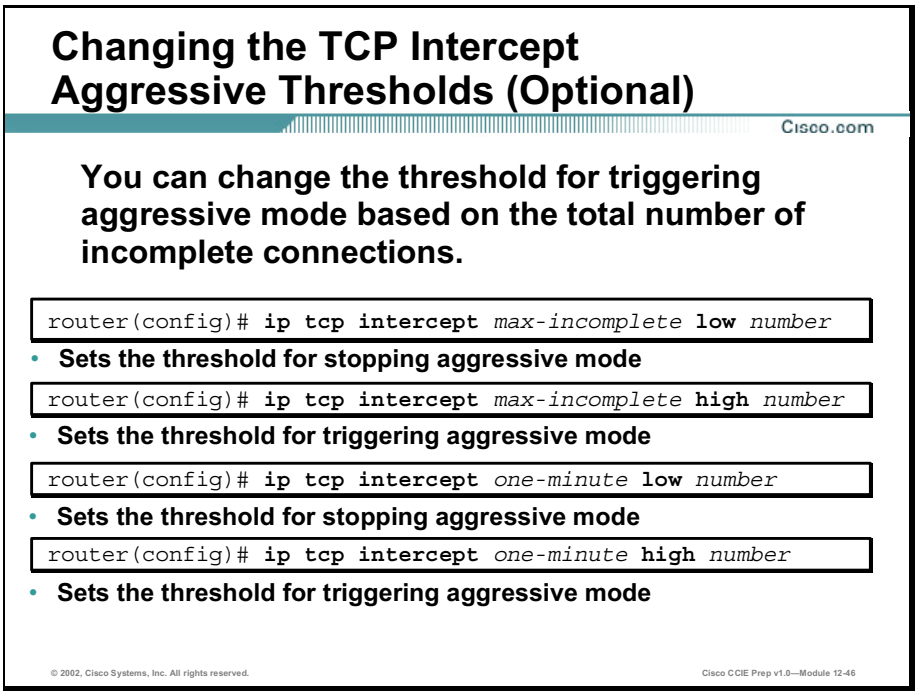

Two factors determine when TCP intercept aggressive behavior begins and ends: total incomplete connections and number of connection requests during the last one-minute sample period. Both of these thresholds have default values that can be redefined.

You can change the threshold for triggering aggressive mode based on the total number of incomplete connections. The default values for **low** and **high** are 900 and 1100 incomplete connections, respectively. To change these values, use the following commands in global configuration mode:

Table 12-23: Change Aggressive Mode Values Based on Incomplete Connections

Step	Command	Description
Step 1	**ip tcp intercept max-incomplete low** *number*	Sets the threshold for stopping aggressive mode
Step 2	**ip tcp intercept max-incomplete high** *number*	Sets the threshold for triggering aggressive mode

You can also change the threshold for triggering aggressive mode based on the number of connection requests received in the last 1-minute sample period. The default values for **low** and **high** are 900 and 1100 connection requests, respectively. To change these values, use the following commands in global configuration mode:

Table 12-24: Change Aggressive Mode Values Based on Connection Requests

Step	Command	Description
Step 1	**ip tcp intercept one-minute low** *number*	Sets the threshold for stopping aggressive mode
Step 2	**ip tcp intercept one-minute high** *number*	Sets the threshold for triggering aggressive mode

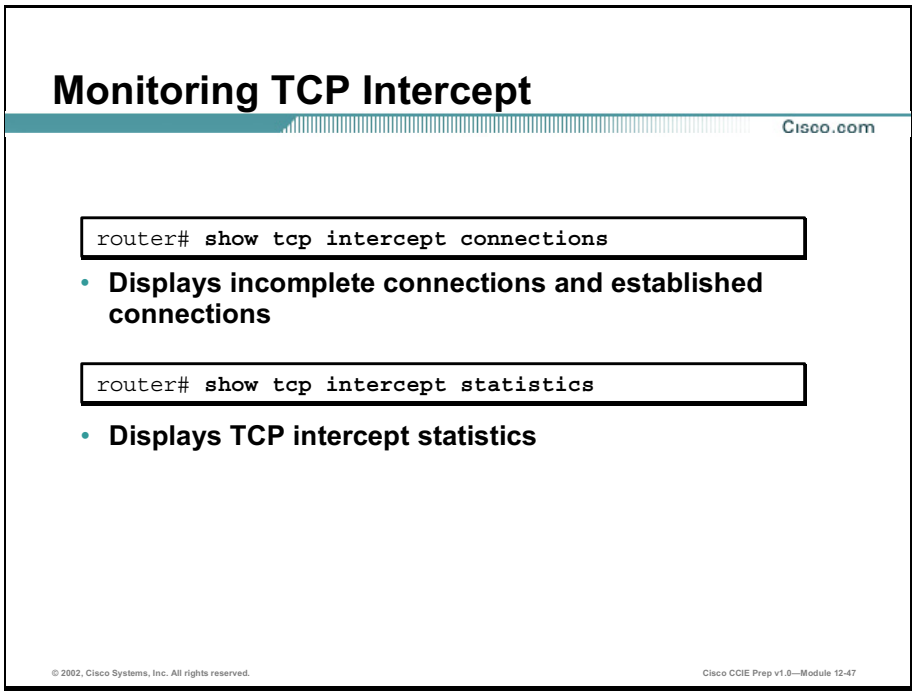

To display TCP intercept information, use the following commands in EXEC mode:

Table 12-25: Display TCP Intercept Information with EXEC Mode Commands

Command	Description
`show tcp intercept connections`	Displays incomplete connections and established connections
`show tcp intercept statistics`	Displays TCP intercept statistics

Context-Based Access Control (CBAC)

This section covers Context-Based Access Control (CBAC).

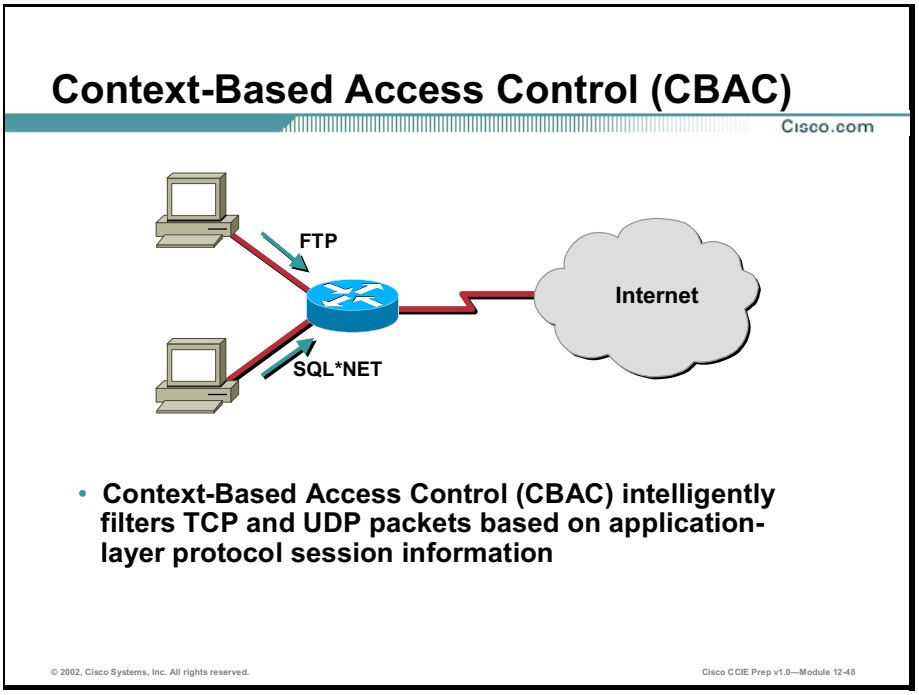

Context-Based Access Control (CBAC) is included in the IOS Firewall Feature Set, now known as Cisco Secure Integrated Software (CSIS). CBAC intelligently filters TCP and User Datagram Protocol (UDP) packets based on application-layer protocol session information. You can configure CBAC to permit specified TCP and UDP traffic through a router only when the connection is initiated from within the internal network you want to protect. CBAC can inspect traffic for sessions that originated from the internal network and allow return traffic from those sessions.

Without CBAC, traffic filtering is limited to access list implementations that examine packets at the network layer, or at most, the transport layer. However, CBAC examines not only network layer and transport layer information but also examines the application-layer protocol information (such as File Transfer Protocol (FTP) connection information) to learn about the state of the session. This allows support of protocols that involve multiple channels created as a result of negotiations in the control channel. Most of the multimedia protocols as well as some other protocols (such as FTP, Remote Procedure Call (RPC), and SQL*Net) involve multiple channels.

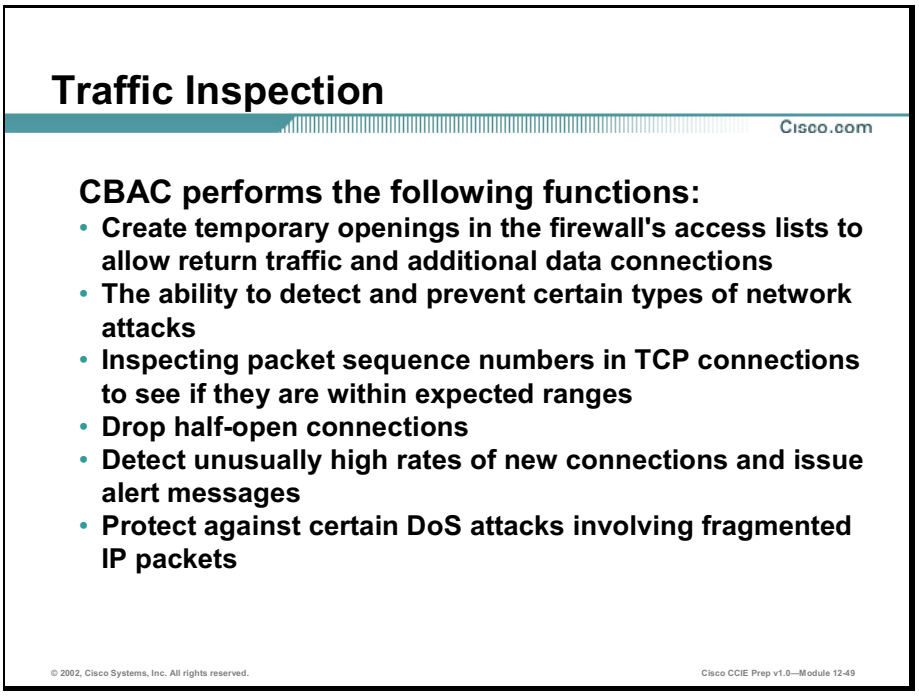

CBAC inspects traffic traveling through the firewall to discover and manage state information for TCP and UDP sessions. The state information is stored locally in memory and is used to perform many functions including:

- Using the state information to create temporary openings in the firewall's access lists to allow return traffic and additional data connections for permissible sessions
- The ability to detect and prevent certain types of network attacks such as SYN-flooding
- Inspecting packet sequence numbers in TCP connections to see if they are within expected ranges—CBAC drops any suspicious packets
- Dropping half-open connections, which require firewall processing and memory resources to maintain
- Detecting unusually high rates of new connections and issuing alert messages
- Protecting against certain DoS attacks involving fragmented IP packets

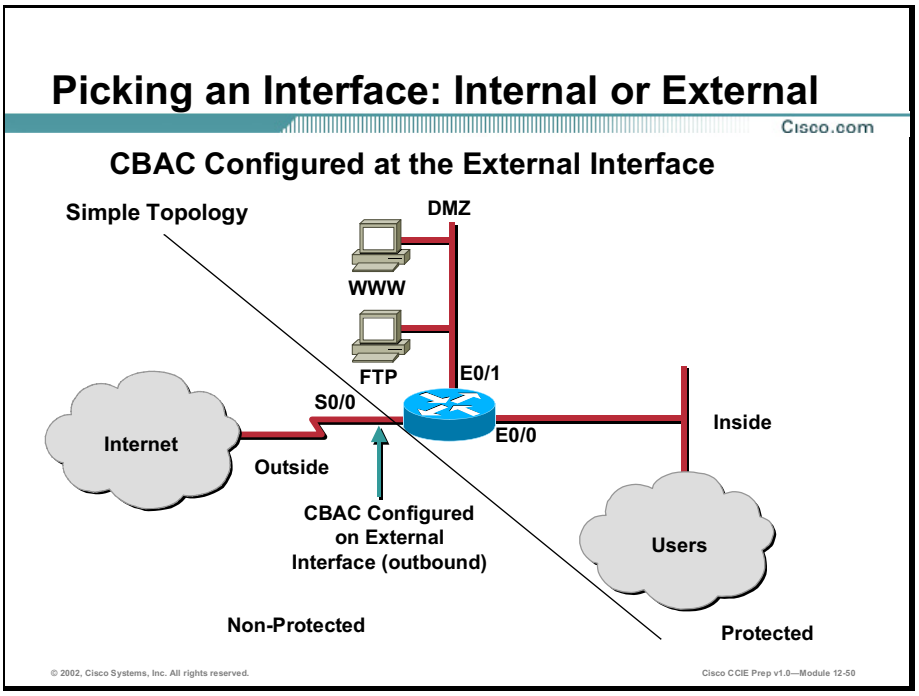

You must decide whether to configure CBAC on an internal or external interface of your firewall.

"Internal" refers to the side where sessions must originate for their traffic to be permitted through the firewall. "External" refers to the side where sessions cannot originate (sessions originating from the external side will be blocked).

The firewall is most commonly used with one of two basic network topologies. Determining which of these topologies is most like your own can help you decide whether to configure CBAC on an internal interface or on an external interface.

This figure shows the first network topology. In this simple topology, CBAC is configured for the external interface Serial 0/0. This prevents specified protocol traffic from entering the firewall, internal network, and Demilitarized Zone (DMZ), unless the traffic is part of a session initiated from within the inspected network.

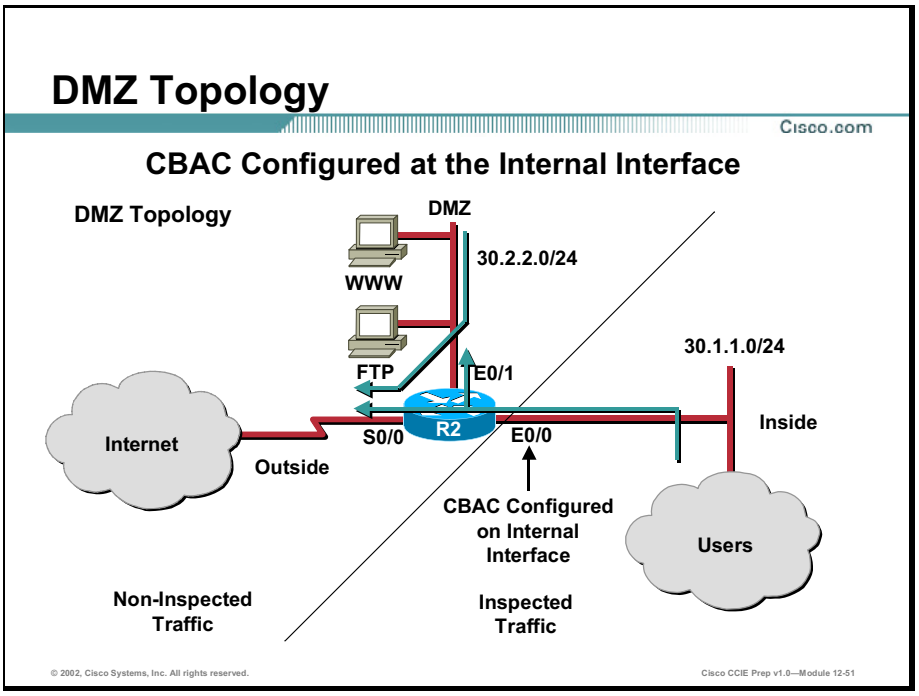

This figure shows the second network topology. In this topology, CBAC is configured for the internal interface Ethernet 0/0. This allows external traffic to access the services in the Demilitarized Zone (DMZ), such as WWW services, but prevents specified protocol traffic from entering your internal network—unless the traffic is part of a session initiated from within the internal network.

Using these two sample topologies, decide whether to configure CBAC on an internal or external interface.

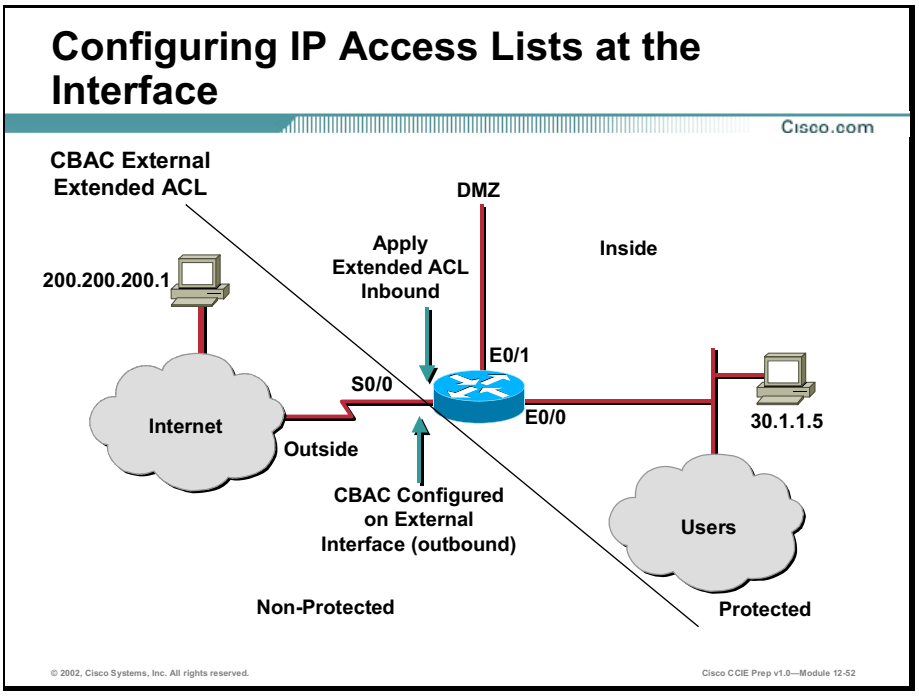

For CBAC to work properly, you need to make sure that you have extended IP access lists configured appropriately at an interface.

Follow these three general rules when evaluating your IP access lists at the firewall:

- **Start with a basic configuration.**

  If you try to configure access lists without a good understanding of how access lists work, you might inadvertently introduce security risks to the firewall and to the protected network. You should be sure you understand what access lists do before you configure your firewall. For more information about access control lists, refer to the "Access Control Lists: Overview and Guidelines" chapter of the Cisco IOS Release 12.0 *Security Configuration Guide*.

  A basic initial configuration allows all network traffic to flow from the protected networks to the unprotected networks, while blocking network traffic from any unprotected networks.

- **Permit CBAC traffic to leave the network through the firewall.**

  All access lists that evaluate traffic leaving the protected network should permit traffic that will be inspected by CBAC. For example, if Telnet will be inspected by CBAC, then Telnet traffic should be permitted on all access lists that apply to traffic leaving the network.

- **Use extended access lists to deny CBAC return traffic entering the network through the firewall.**

  For temporary openings to be created in an access list, the access list must be an extended access list. So wherever you have access lists that will be applied to returning traffic, you

must use extended access lists. The access lists should deny CBAC return traffic because CBAC will open up temporary holes in the access lists. (You want traffic to be normally blocked when it enters your network.)

Note	If your firewall only has two connections, one to the internal network and one to the external network, using all inbound access lists works well because packets are stopped before they get a chance to affect the router itself.

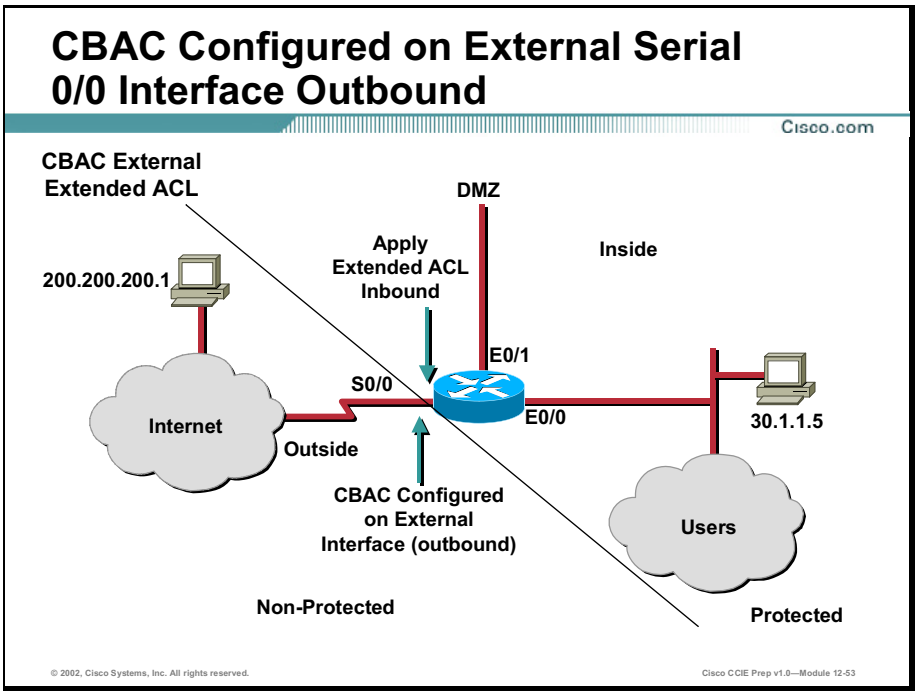

In this scenario, CBAC has been configured on the external Serial 0/0 interface outbound. This means any traffic leaving the inside or DMZ network will have its session state information maintained in the router. Dynamic Access Control List (ACL) entries will be applied to returning (inbound) traffic to allow this specific traffic inside to the protected network(s).

These dynamic entries will be applied to an inbound extended access list, so also on interface serial 0/0 you create an extended access list and apply it inbound to this interface. This extended access list can be as simple as:

```
access-list 100 deny ip any any
```

This will block any traffic initiated from the outside heading into the router.

When a session is started from the inside heading toward the Internet, CBAC will create a dynamic entry and apply it to the inbound access list. For instance, say a client on the inside at IP address 30.1.1.5 requested a web page on the Internet located at 200.200.200.1.

The following steps will occur:

Step 1 Traffic leaves the inside network, travels through the router, and heads toward interface serial 0/0.

Step 2 Serial 0/0 has an inspection rule stipulating that this traffic should be CBAC inspected.

Step 3 A dynamic entry is applied to extended access list 100 to allow this return traffic. At this time the extended access list would look something like this:

```
access-list 100 permit tcp host 200.200.200.1 eq 80 host 30.1.1.5 eq 1044
access-list 100 deny ip any any
```

Step 4 The packet is forwarded into the Internet.

Step 5 Return traffic can now safely pass the extended access list and head into the inside network.

Step 6 After a configurable amount of time when no traffic has passed, the dynamic entry will be removed from extended access list 100.

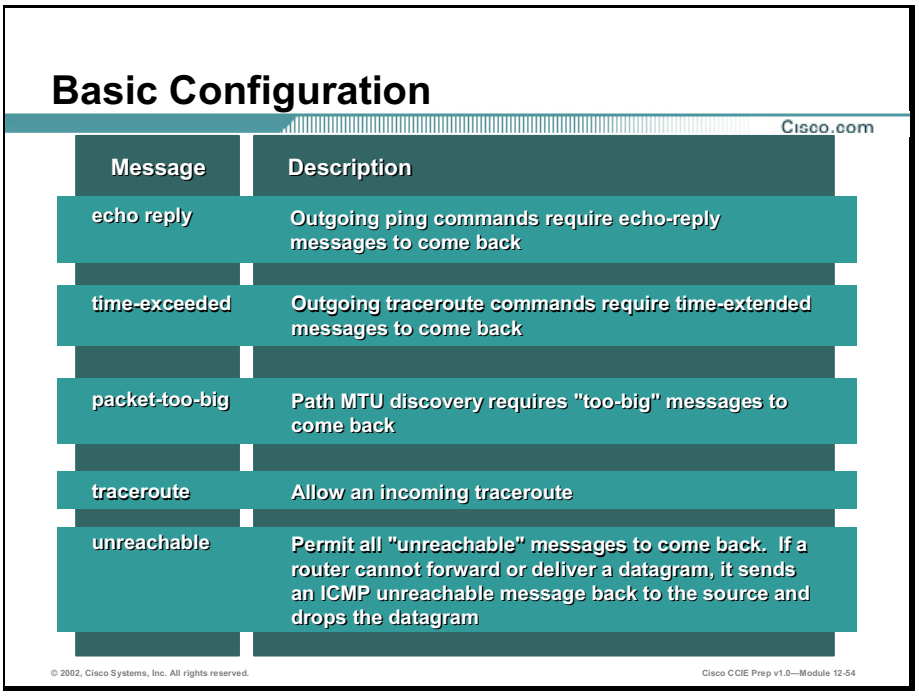

The first time you configure the Cisco IOS firewall, it is helpful to start with a basic access list configuration that makes the operation of the firewall easy to understand without compromising security. The basic configuration allows all network traffic from the protected networks access to the unprotected networks, while blocking all network traffic (with some exceptions) from the unprotected networks to the protected networks.

Any firewall configuration depends on your site security policy. If the basic configuration does not meet your initial site security requirements, configure the firewall to meet your policy.

For example, you might want Internet Control Message Protocol (ICMP) ping and Traceroute traffic to pass into your firewall. To do that, you modify your extended access list to allow this traffic through. Why? Because ICMP traffic is not inspected by CBAC, meaning specific entries are needed in the access list to permit return traffic for ICMP commands. For example, a user on a protected network uses the **ping** command to get the status of a host on an unprotected network; without entries in the access list that permit **echo reply** messages, the user on the protected network gets no response to the **ping** command.

The following table lists entries you might configure to permit certain ICMP messages.

Table 12-26: ICMP Messages

Message	Description
echo reply	Outgoing ping commands require echo-reply messages to come back.
time-exceeded	Outgoing traceroute commands require time-exceeded messages to come back.
packet-too-big	Path MTU discovery requires "too-big" messages to come back.
traceroute	Allows an incoming traceroute.
unreachable	Permits all "unreachable" messages to come back. If a router cannot forward or deliver a datagram, it sends an ICMP unreachable message back to the source and drops the datagram.

At this point, your extended access list would look something like this:

```
access-list 100 permit icmp any any eq echo-reply
access-list 100 permit icmp any any eq time-exceeded
access-list 100 deny ip any any
```

You might also want to implement anti-spoofing protection. For example, your Inside network is 30.1.1.0/24. You can safely assume that no packet should be generated from the Internet using this network, so you add the following:

```
access-list 100 deny ip 30.1.1.0 0.0.0.255 any
```

You might also want to prevent broadcast attacks. To do so, use the following entry:

```
access-list 100 deny ip host 255.255.255.255 any
```

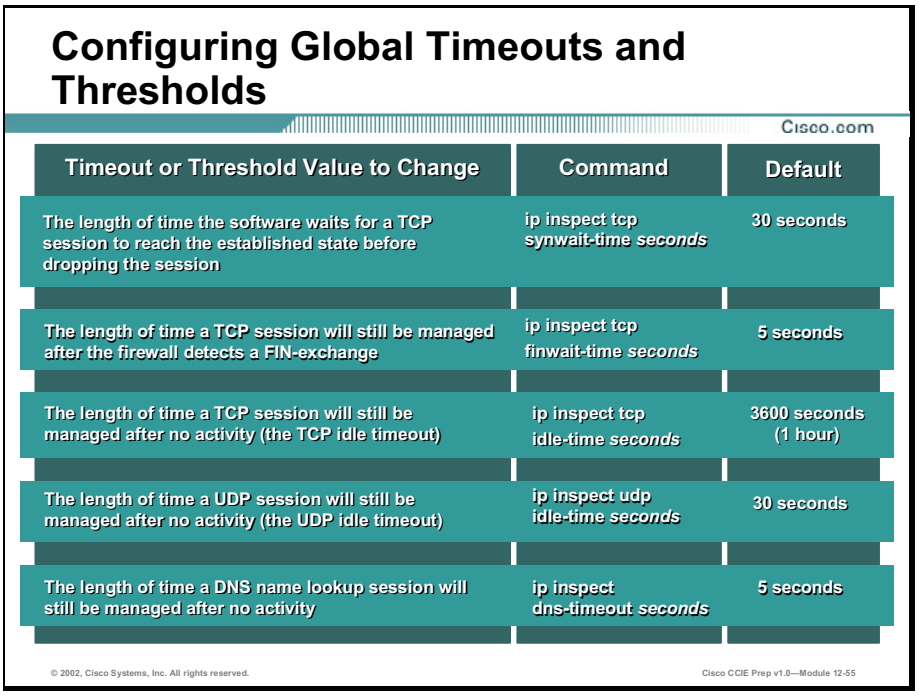

CBAC uses timeouts and thresholds to determine how long to manage state information for a session, and to determine when to drop sessions that do not become fully established. These timeouts and thresholds apply globally to all sessions.

You can use the default timeout and threshold values, or you can change to values more suitable to your security requirements. You should make any changes to the timeout and threshold values before you continue configuring CBAC.

All the available CBAC timeouts and thresholds are listed in the table shown along with the corresponding command and default value.

Configuring Global Timeouts and Thresholds (Cont.)

Cisco.com

Timeout or Threshold Value to Change	Command	Default
The number of existing half-open sessions that will cause the software to start deleting half-open sessions	ip inspect max-incomplete high *number*	500 existing half-open sessions
The rate of new sessions that will cause the software to start deleting half-open sessions	ip inspect one-minute high *number*	500 half-open sessions per minute
The rate of new sessions that will cause the software to stop deleting half-open sessions	ip inspect one-minute low *number*	400 half-open sessions per minute
The number of existing half-open TCP sessions with the same destination host address that will cause the software to start dropping half-open sessions to the same destination host address	ip inspect tcp max-incomplete host *number* block-time *minutes*	50 existing half-open TCP sessions; 0 minutes

To change a global timeout or threshold listed in the "Timeout or Threshold Value to Change" column, use the global configuration command in the "Command" column:

Table 12-27: Timeout and Threshold Value Commands

Timeout and Threshold Value to Change	Command	Default
The length of time the software waits for a TCP session to reach the established state before dropping the session	**ip inspect tcp synwait-time** *seconds*	30 seconds
The length of time a TCP session will still be managed after the firewall detects a FIN-exchange	**ip inspect tcp finwait-time** *seconds*	5 seconds
The length of time a TCP session will still be managed after no activity (the TCP idle timeout)	**ip inspect tcp idle-time** *seconds*	3600 seconds (1 hour)
The length of time a UDP session will still be managed after no activity (the UDP idle timeout)	**ip inspect udp idle-time** *seconds*	30 seconds
The length of time a Domain Name Service (DNS) name lookup session will still be managed after no activity	**ip inspect dns-timeout** *seconds*	5 seconds
The number of existing half-open sessions that will cause the software to start deleting half-open sessions	**ip inspect max-incomplete high** *number*	500 existing half-open sessions
The rate of new sessions that will cause the software to start deleting half-open sessions	**ip inspect one-minute high** *number*	500 half-open sessions per minute
The rate of new sessions that will cause the software to stop deleting half-open sessions	**ip inspect one-minute low** *number*	400 half-open sessions per minute
The number of existing half-open TCP sessions with the same destination host address that will cause the software to start dropping half-open sessions to the same destination host address	**ip inspect tcp max-incomplete host** *number* **block-time** *minutes*	50 existing half-open TCP sessions; 0 minutes

To reset any threshold or timeout to the default value, use the **no** form of the command in the table shown.

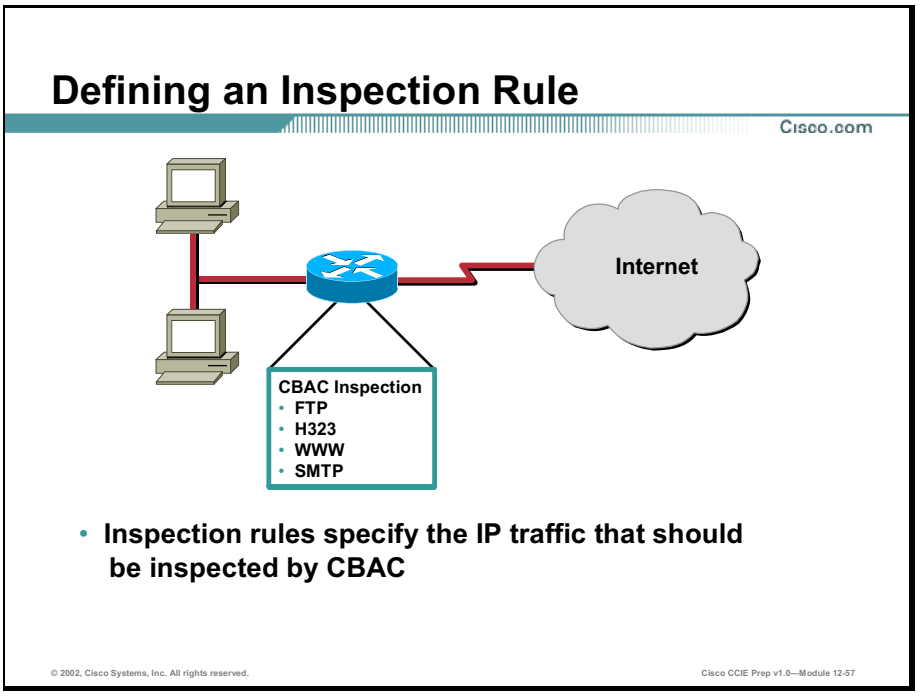

After you configure global timeouts and thresholds, you must define an inspection rule. This rule specifies what IP traffic (which application-layer protocols) will be inspected by CBAC at an interface.

Normally, you define only one inspection rule.

An inspection rule should specify each desired application-layer protocol as well as generic TCP or generic UDP if desired. The inspection rule consists of a series of statements, each listing a protocol, and specifying the same inspection rule name.

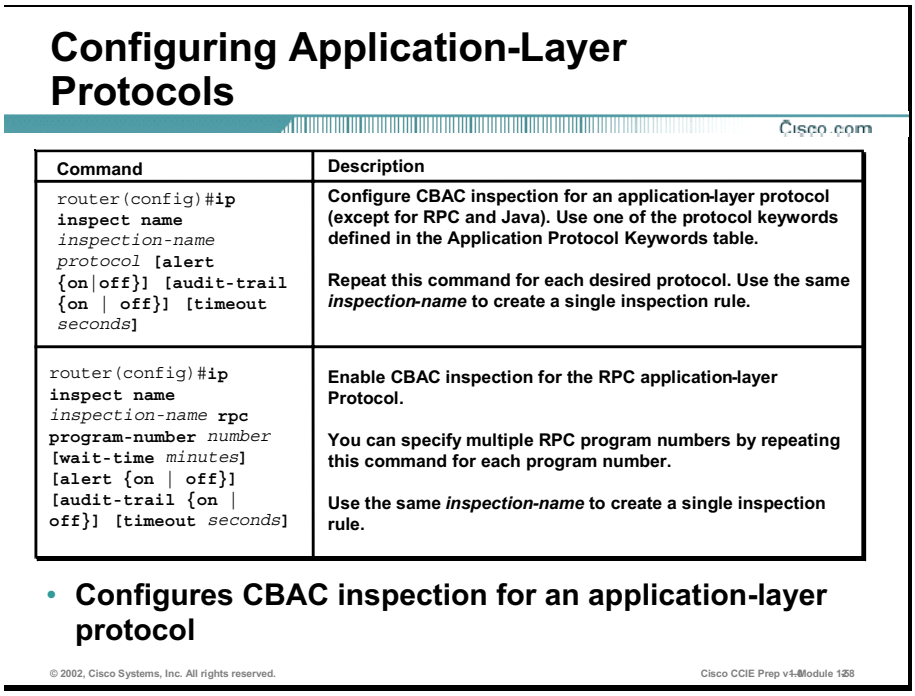

To configure CBAC inspection for an application-layer protocol, use one or both of the following global configuration commands:

Table 12-28: CBAC Inspection Configuration Commands

Command	Description
`router(config)# ip inspect name inspection-name protocol [alert {on|off}] [audit-trail {on | off}] [timeout seconds]`	Configure CBAC inspection for an application-layer protocol (except for RPC and Java). Use one of the protocol keywords defined in the Application Protocol Keywords table. Repeat this command for each desired protocol. Use the same *inspection-name* to create a single inspection rule.
`router(config)# ip inspect name inspection-name rpc program-number number [wait-time minutes] [alert {on | off}] [audit-trail {on | off}] [timeout seconds]`	Enable CBAC inspection for the RPC application-layer protocol. You can specify multiple RPC program numbers by repeating this command for each program number. Use the same *inspection-name* to create a single inspection rule.

The following table identifies application protocol keywords.

Table 12-29: Application Protocol Keywords

Application Protocol	*protocol* Keyword
CU-SeeMe	**cuseeme**
FTP	**ftp**
H.323	**h323**
Microsoft NetShow	**netshow**
UNIX R commands (rlogin, rexec, rsh)	**rcmd**
RealAudio	**realaudio**
SMTP	**smtp**
RPC	**rpc**
SQL*Net	**sqlnet**
StreamWorks	**streamworks**
TFTP	**tftp**
VDOLive	**vdolive**

Configuring Java Inspection

Cisco.com

Step	Command	Description
1.	router(config)#**ip access-list standard** *name* **(Use permit and deny statements as appropriate.)** **permit ...** **deny ...** **or** router(config)#**access-list** *access-list-number* **{deny | permit}** *source* **[***source-wildcard***]**	**Creates a standard access list that permits traffic only from friendly sites, and denies traffic from hostile sites.** **If you want all internal users to be able to download friendly applets, use the any keyword for the destination as appropriate---but be careful to not misuse the any keyword to inadvertently allow all applets through.**
2.	router(config)#**ip inspect name inspection-name http [java-list** *access-list***] [alert {on | off}] [audit-trail {on | off}] [timeout** *seconds***]**	**Blocks all Java applets except for applets from the friendly sites defined previously in the access-list. Java blocking only works with standard access-lists.** **Use the same *inspection-name* as when you specified other protocols, to create a single inspection rule.**

- **Blocks all Java applets except for applets from friendly locations**

Java applet filtering distinguishes between trusted and untrusted applets by relying on a list of external sites that you designate as "friendly." If an applet is from a friendly site, the firewall allows the applet through. If the applet is not from a friendly site, the applet will be blocked. (Alternatively, you could permit applets from all external sites except for those you specifically designate as hostile.)

To block all Java applets except for applets from friendly locations, use the following global configuration commands:

Table 12-30: Block Java Applets Commands

Step	Command	Description
Step 1	router(config)# **ip access-list standard** *name* **permit** ... **deny** ... (Use permit and deny statements as appropriate.) or router(config)# **access-list** *access-list-number* {**deny** | **permit**} *source [source-wildcard]*	Creates a standard access list that permits traffic only from friendly sites, and denies traffic from hostile sites If you want all internal users to be able to download friendly applets, use the **any** keyword for the destination as appropriate---but be careful to not misuse the **any** keyword to inadvertently allow all applets through.
Step 2	router(config)# **ip inspect name** *inspection-name* **http** [**java-list** *access-list*] [**alert** {**on** | **off**}] [**audit-trail** {**on** | **off**}] [**timeout** *seconds*]	Blocks all Java applets except for applets from the friendly sites defined previously in the access list. Java blocking only works with standard access lists. Use the same *inspection-name* as when you specified other protocols, to create a single inspection rule.

To configure CBAC inspection rules for IP fragmentation checking, use the following form of the **ip inspect name** global configuration command:

Table 12-31: ip inspect name Command

Command	Description
router(config)# **ip inspect name** *inspection-name* **fragment** [**max** *number* timeout ***number***]	Configures IP fragmentation checking in CBAC inspection rules

Repeat this command for each named inspection rule in which you want to inspect IP fragments.

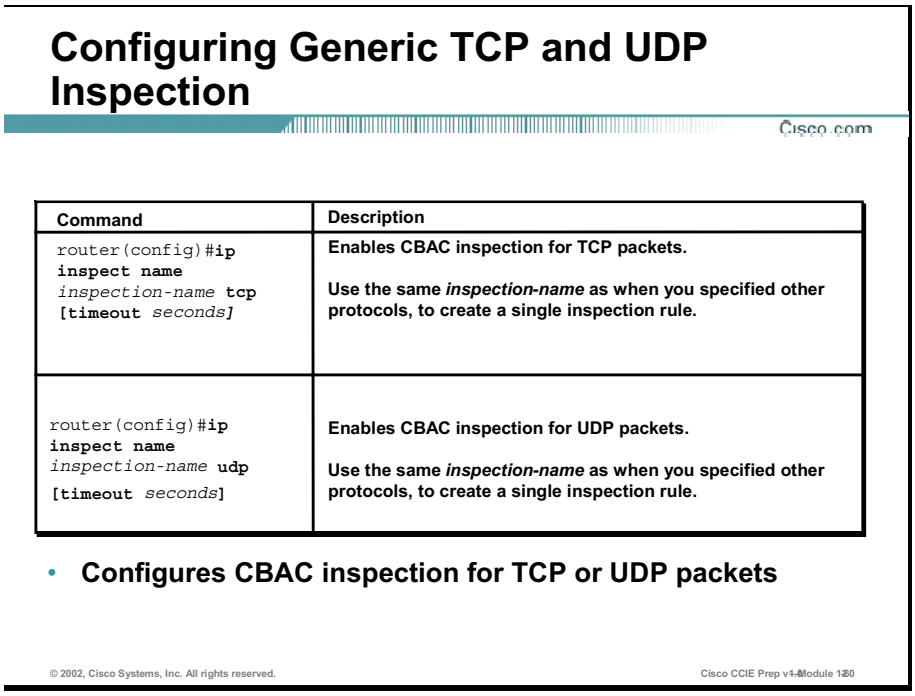

You can configure TCP and UDP inspection to permit TCP and UDP packets to enter the internal network through the firewall, even if the application-layer protocol is not configured for inspection. However, TCP and UDP inspection do not recognize application-specific commands, and therefore might not permit all return packets for an application, particularly if the return packets have a different port number than the previous exiting packet.

To configure CBAC inspection for TCP or UDP packets, use one or both of the following global configuration commands:

Table 12-32: CBAC Inspection Commands

Command	Description
router(config)# **ip inspect name** *inspection-name* **tcp** [**timeout** *seconds*]	Enables CBAC inspection for TCP packets. Use the same *inspection-name* as when you specified other protocols, to create a single inspection rule.
router(config)# **ip inspect name** *inspection-name* **udp** [**timeout** *seconds*]	Enables CBAC inspection for UDP packets. Use the same *inspection-name* as when you specified other protocols, to create a single inspection rule.

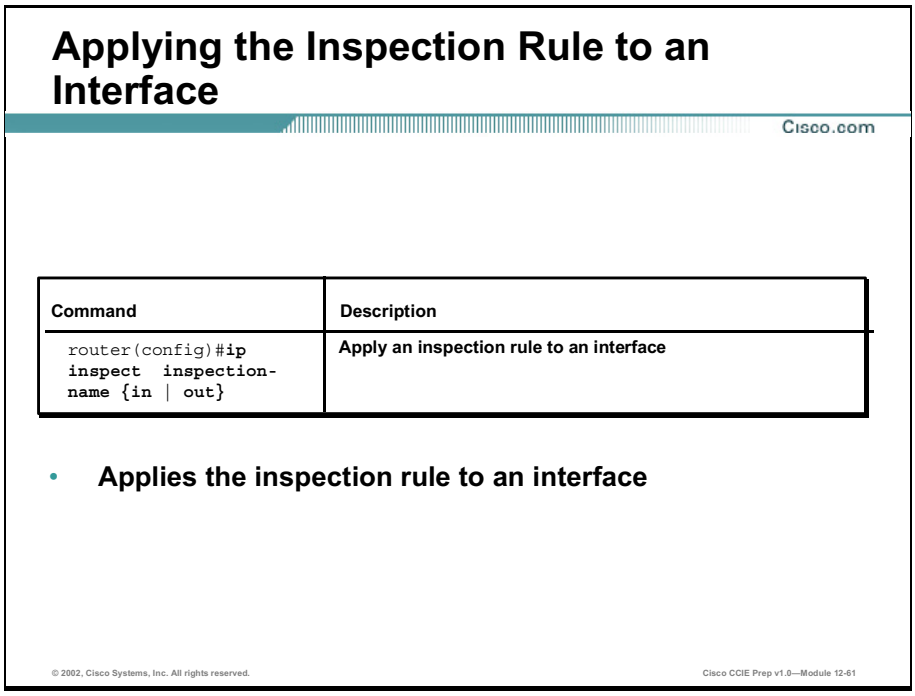

After you define an inspection rule, you apply that rule to an interface.

Normally, you apply only one inspection rule to one interface.

To apply an inspection rule to an interface, use the following interface configuration command:

Table 12-33: < router (config-if)**# ip inspect** inspection-name {**in** | **out**}**> Command**

Command	Description
router(config-if)**# ip inspect** inspection-name {**in** | **out**}	Apply an inspection rule to an interface

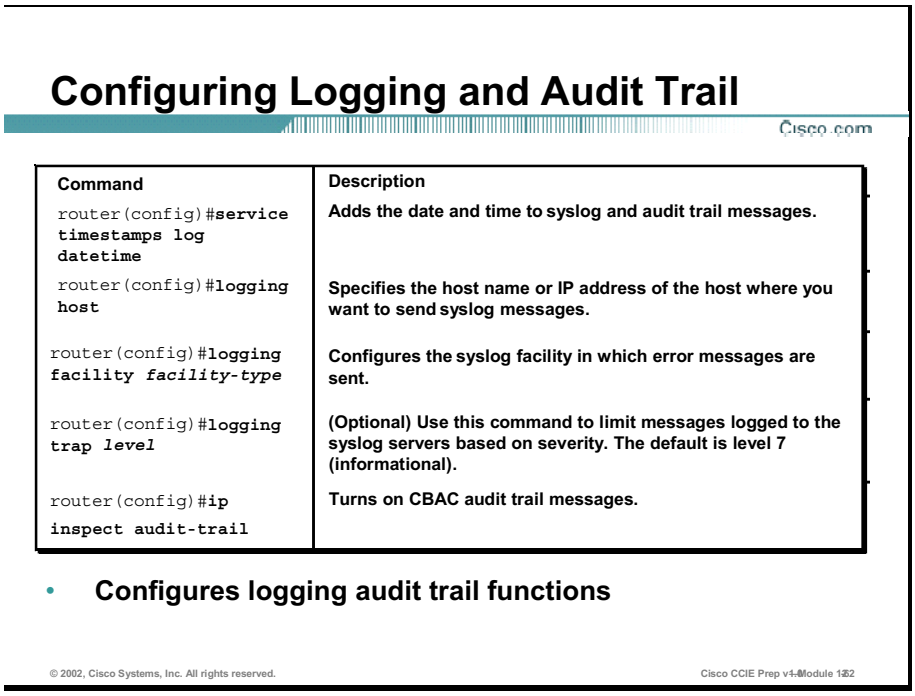

Turn on logging and audit trail to provide a record of network access through the firewall, including illegitimate access attempts, and inbound and outbound services. To configure logging and audit trail functions, enter the following commands in global configuration mode:

Table 12-34: Logging and Audit Trail Commands

Command	Description
`router(config)# service timestamps log datetime`	Adds the date and time to syslog and audit trail messages.
`router(config)# logging host`	Specifies the host name or IP address of the host where you want to send syslog messages.
`router(config)# logging facility facility-type`	Configures the syslog facility in which error messages are sent.
`router(config)# logging trap level`	(Optional) Use this command to limit messages logged to the syslog servers based on severity. The default is level 7 (informational).
`router(config)# ip inspect audit-trail`	Turns on CBAC audit trail messages.

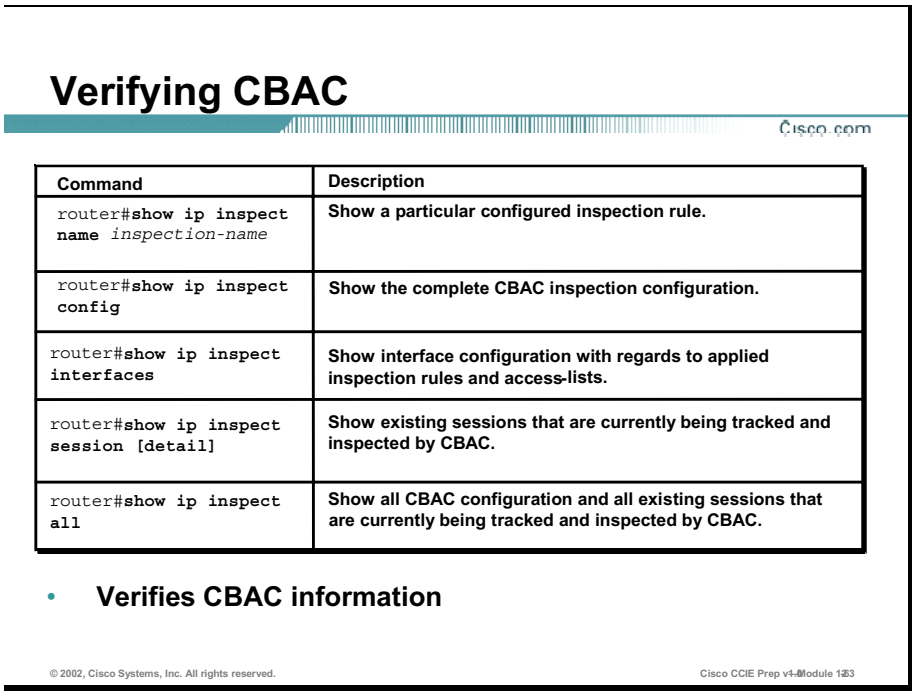

You can verify CBAC information by using one or more of the following EXEC commands:

Table 12-35: Verify CBAC Information with EXEC Commands

Command	Description
router# **show ip inspect name** *inspection-name*	Show a particular configured inspection rule.
router# **show ip inspect config**	Show the complete CBAC inspection configuration.
router# **show ip inspect interfaces**	Show interface configuration with regards to applied inspection rules and access lists.
router# **show ip inspect session** [**detail**]	Show existing sessions that are currently being tracked and inspected by CBAC.
router# **show ip inspect all**	Show all CBAC configuration and all existing sessions that are currently being tracked and inspected by CBAC.

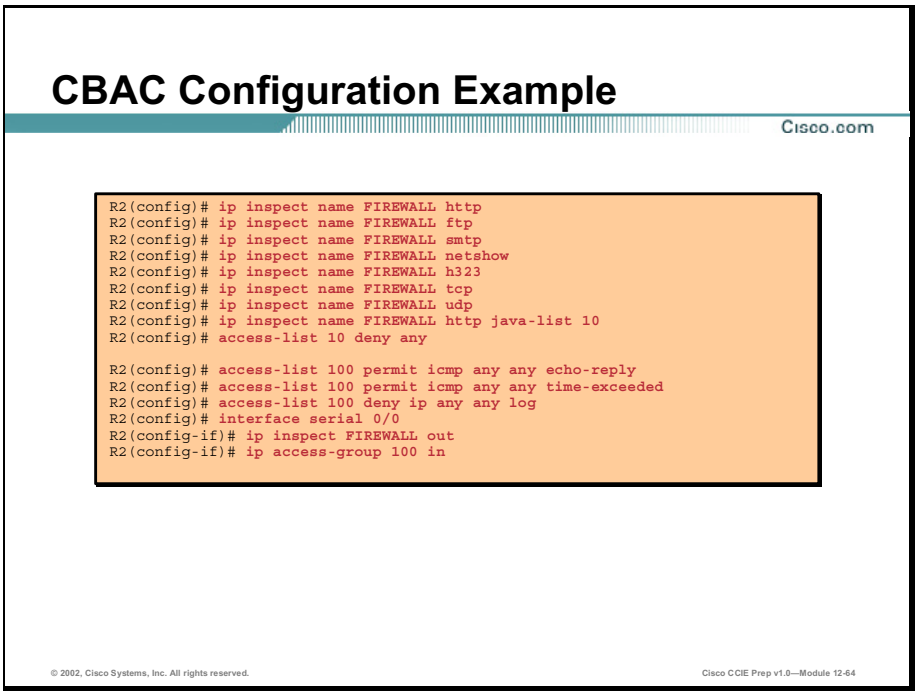

Now take a look a putting it all together.

In this scenario, you will configure CBAC on the external interface S0/0, to inspect outbound traffic. This will allow traffic from clients on the inside to receive their return traffic from the outside. You will configure an inbound extended access list on the external interface, S0/0, to allow certain ICMP traffic into the protected network and block all other traffic.

```
R2(config)# ip inspect name FIREWALL http
R2(config)# ip inspect name FIREWALL ftp
R2(config)# ip inspect name FIREWALL smtp
R2(config)# ip inspect name FIREWALL netshow
R2(config)# ip inspect name FIREWALL h323
R2(config)# ip inspect name FIREWALL tcp
R2(config)# ip inspect name FIREWALL udp
R2(config)# ip inspect name FIREWALL http java-list 10
R2(config)# access-list 10 deny any
R2(config)# access-list 100 permit icmp any any echo-reply
R2(config)# access-list 100 permit icmp any any time-exceeded
R2(config)# access-list 100 deny ip any any log
R2(config)# interface Serial 0/0
R2(config-if)# ip access-group 100 in
R2(config-if)# ip inspect FIREWALL out
```

- The DMZ network can reach the Outside network.
- Outside and DMZ cannot initiate a ping into the Inside network, but can reply to pings.

- The Inside network can traceroute to any Outside or DMZ host.
- The Inside network can reach the Outside network, and the dynamic entries created on access list 100 will allow the return traffic.

IPSec

This section covers the creation of Internet Protocol Security (IPSec) tunnels in order to pass data from R3 to R4 over an insecure Internet connection.

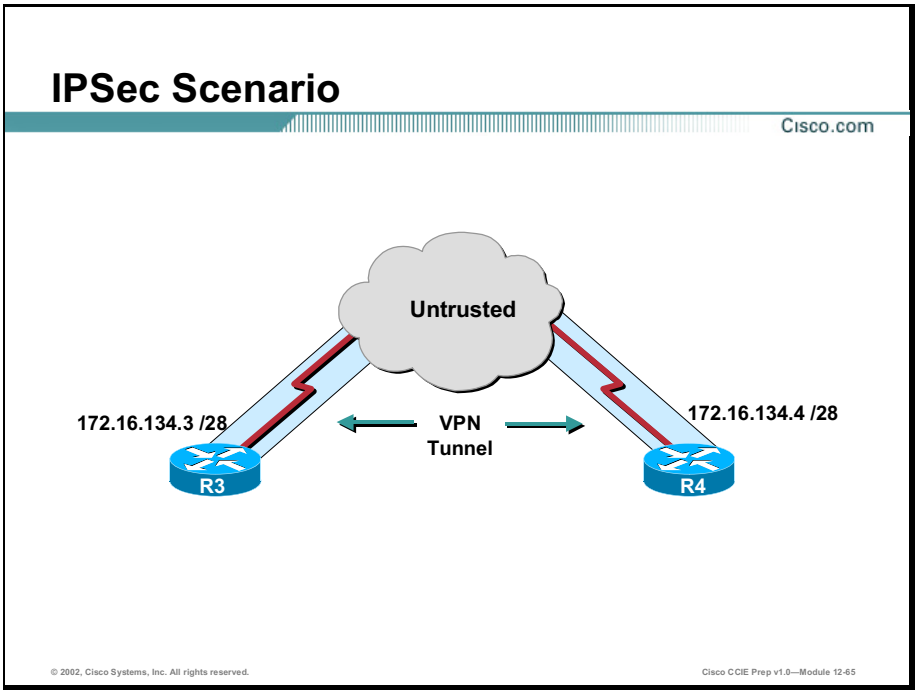

In this scenario, you will create an IPSec tunnel that will allow you to pass data from R3 to R4. In this process, you will encrypt the traffic.

These are the steps you need to take in order to configure a secure Virtual Private Network (VPN) tunnel.

```
R3(config)# crypto isakmp policy 10
R3(config-isakmp)# hash md5
R3(config-isakmp)# authentication pre-share
R3(config-isakmp)# group 2
R3(config-isakmp)# exit
R3(config)# crypto isakmp key ciscoCCIE address 172.16.134.4
R3(config)# crypto ipsec transform-set TSET esp-des esp-md5-hmac
R3(cfg-crypto-trans)# exit
R3(config)# access-list 101 permit ip host 172.16.134.3 host 172.16.134.4
R3(config)# crypto map MYMAP 10 ipsec-isakmp
R3(config-crypto-map)# set peer 172.16.134.4
R3(config-crypto-map)# set transform-set TSET
R3(config-crypto-map)# match address 101
R3(config-crypto-map)# exit
R3(config)# int serial1/0
```

```
R3(config-if)# crypto map MYMAP
```

Now that you have the configuration, you can now delve into the details of this configuration.

First, define an Internet Security Association and Key Management Protocol (ISAKMP) policy that will use Message Digest Version 5 (md5) hashing, Diffie-Hellman group 2, and pre-shared key for authentication.

```
R3(config)# crypto isakmp policy 10
R3(config-isakmp)# hash md5
R3(config-isakmp)# authentication pre-share
R3(config-isakmp)# group 2
```

Now that you have defined ISAKMP policy, set up the pre-shared key that will be used for authentication. You must remember this pre-shared key is case-sensitive and needs to be the same at both ends of the tunnel. In this command, you are defining the pre-shared key and also defining the IP address of the peer (R4).

```
R3(config)# crypto isakmp key ciscoCCIE address 172.16.134.4
```

At this point, your ISAKMP configuration is pretty much done. Now you need to setup a transform set that will encrypt all the data that will be flowing through the tunnel. One important thing to remember is the fact that you can setup multiple transform sets, but at least one transform set should match at both ends for any kind of security association to take place.

```
R3(config)# crypto ipsec transform-set TSET esp-des esp-md5-hmac
```

Now set up an access list that will define traffic that is allowed across this tunnel. Here, you are allowing all IP traffic to flow through the encrypted tunnel from host 172.16.134.3 (R3) to 172.16.143.4 (R4).

```
R3(config)# access-list 101 permit ip host 172.16.134.3 host 172.16.134.4
```

Once the access list has been created, set up a crypto map that sets the IPSec peer to 172.16.134.4 (R4) and points to the transform-set called Transform-s+B140et (TSET). The match address command is used to specify an extended access list for the crypto map entry.

```
R3(config)# crypto map MYMAP 10 ipsec-isakmp
R3(config-crypto-map)# set peer 172.16.134.4
R3(config-crypto-map)# set transform-set TSET
R3(config-crypto-map)# match address 101
```

The crypto map has been set up and now you need to apply this map to the serial 1 interface of R3. Serial 1/0 is the interface that is part of the Internet cloud.

```
R3(config)# int s1/0
R3(config-if)# crypto map MYMAP
```

Remember that you will be required to do similar configuration on R4 for this to work. Here is the configuration that needs to be done on R4.

```
R4(config)# crypto isakmp policy 10
R4(config-isakmp)# hash md5
R4(config-isakmp)# authentication pre-share
R4(config-isakmp)# group 2
R4(config-isakmp)# exit
R4(config)# crypto isakmp key ciscoCCIE address 172.16.134.3
R4(config)# crypto ipsec transform-set TSET esp-des esp-md5-hmac
R4(cfg-crypto-trans)# exit
R4(config)# access-list 101 permit ip host 172.16.134.4 host 172.16.134.3
R4(config)# crypto map MYMAP 10 ipsec-isakmp
R4(config-crypto-map)# set peer 172.16.134.3
R4(config-crypto-map)# set transform-set TSET
R4(config-crypto-map)# match address 101
R4(config-crypto-map)# exit
R4(config)# int s0/1
R4(config-if)# crypto map MYMAP
R4(config-if)# exit
```

Once the configuration is done, you will use the following commands to verify your configuration.

```
R4# ping
Protocol [ip]:
Target IP address: 172.16.134.3
Repeat count [5]: 20
Datagram size [100]:
Timeout in seconds [2]:
Extended commands [n]: y
Source address or interface: 172.16.134.4
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 20, 100-byte ICMP Echos to 172.16.134.3, timeout is 2 seconds:
..!!!!!!!!!!!!!!!!!!
Success rate is 100 percent (18/20), round-trip min/avg/max = 20/22/24 ms

R4# show crypto map
Crypto Map "MYMAP" 10 ipsec-isakmp
        Peer = 172.16.134.3
        Extended IP access list 101
            access-list 101 permit ip host 172.16.134.4 host 172.16.134.3
        Current peer: 172.16.134.3
        Security association lifetime: 4608000 kilobytes/3600 seconds
        PFS (Y/N): N
        Transform sets={ TSET, }
        Interfaces using crypto map MYMAP:
                Serial0/1

R4# show crypto ipsec sa

interface: Serial0/1
    Crypto map tag: MYMAP, local addr. 172.16.134.4

   local  ident (addr/mask/prot/port):  (172.16.134.4/255.255.255.255/0/0)
   remote ident (addr/mask/prot/port): (172.16.134.3/255.255.255.255/0/0)
   current_peer: 172.16.134.3
     PERMIT, flags={origin_is_acl,}
    #pkts encaps: 39, #pkts encrypt: 39, #pkts digest 39
    #pkts decaps: 39, #pkts decrypt: 39, #pkts verify 39
    #pkts compressed: 0, #pkts decompressed: 0
    #pkts not compressed: 0, #pkts compr. failed: 0, #pkts decompress failed: 0
    #send errors 0, #recv errors 0

     local crypto endpt.: 172.16.134.4 , remote crypto endpt.: 172.16.134.3
     path mtu 1500, media mtu 1500
     current outbound spi: 1AD32448

     inbound esp sas:
      spi: 0x1BD4030D(466879245)
```

```
    transform: esp-des esp-md5-hmac ,
    in use settings ={Tunnel, }
    slot: 0, conn id: 2000, flow_id: 1, crypto map: MYMAP
    sa timing: remaining key lifetime (k/sec): (4607994/3461)
    IV size: 8 bytes
    replay detection support: Y

inbound ah sas:

inbound pcp sas:

outbound esp sas:
 spi: 0x1AD32448(450045000)
    transform: esp-des esp-md5-hmac ,
    in use settings ={Tunnel, }
    slot: 0, conn id: 2001, flow_id: 2, crypto map: MYMAP
    sa timing: remaining key lifetime (k/sec): (4607994/3452)
    IV size: 8 bytes
    replay detection support: Y

outbound ah sas:

outbound pcp sas:
```

Summary

This section summarizes the key points discussed in this lesson.

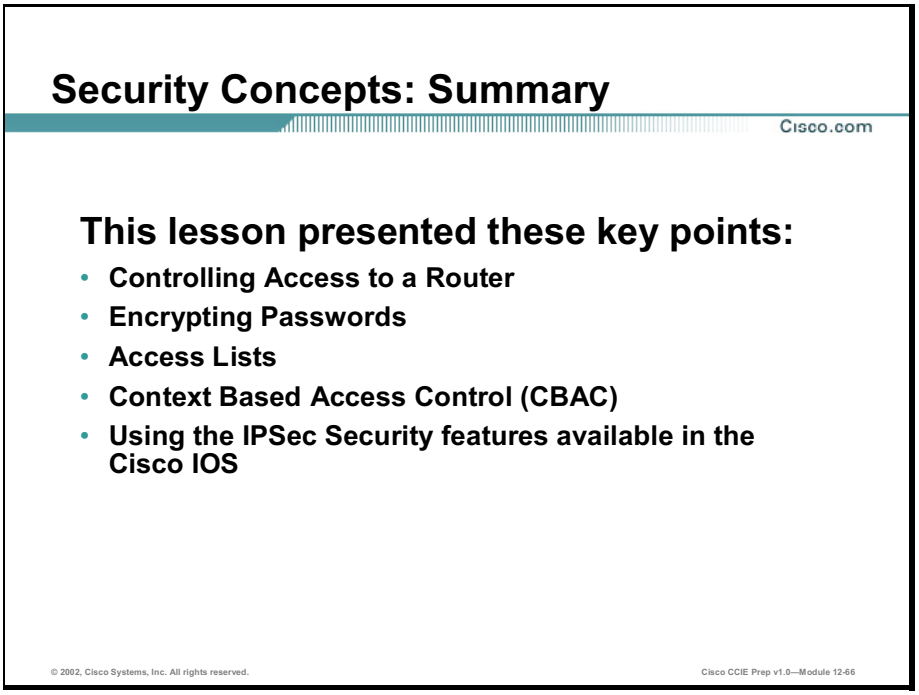

Next Steps

After completing this lesson, go to:

- VoIP Concepts

References

For additional information, refer to these resources:

- http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121cgcr/secur_c/index.htm

Lesson Assessment (Quiz)

Q1) Passwords can be assigned to which of the following lines?

A) Aux

B) TTY

C) VTY

D) All of the above

Q2) What command would you enter on the VTY lines to allow a user Telneting into the router direct access to privilege mode without entering the enable password?

Q3) Which type of access list is used to implement Lock and Key?

A) Named

B) Time-based

C) Dynamic

D) Reflexive

Q4) Which are the two TCP Intercept modes supported on an IOS router?

A) Reset

B) Intercept

C) Watch

D) Block

Q5) Which of the following are used to encrypt data between two routers?

A) CBAC

B) IPSec

C) TCP Interface

D) GRE

Voice over IP Concepts

Overview

Voice over IP (VoIP) is a technology for call setup, transmission, and call teardown over an Internet Protocol (IP) network. This lesson will teach the basics of VoIP as they relate to the Cisco Certified Internetwork Expert (CCIE) lab.

Importance

VoIP equipment is part of the CCIE lab candidate's rack and, therefore, should be an expected lab objective.

Objectives

Upon completing this lesson, you will be able to:

- Perform VoIP configuration
- Configure advanced VoIP features

Learner Skills and Knowledge

To fully benefit from this lesson, you must have these prerequisite skills and knowledge:

- Passed the CCIE written qualification exam
- Completed the Cisco Voice over Frame Relay, ATM, and IP (CVOICE) course or have the equivalent knowledge

Outline

This lesson includes these sections:

- Overview
- VoIP Configuration
- Advanced VoIP Features
- Summary
- Lesson Assessment (Quiz)

Voice over IP Configuration

This section details the configuration of Plain Old Telephone Service (POTS) and VoIP dial-peers.

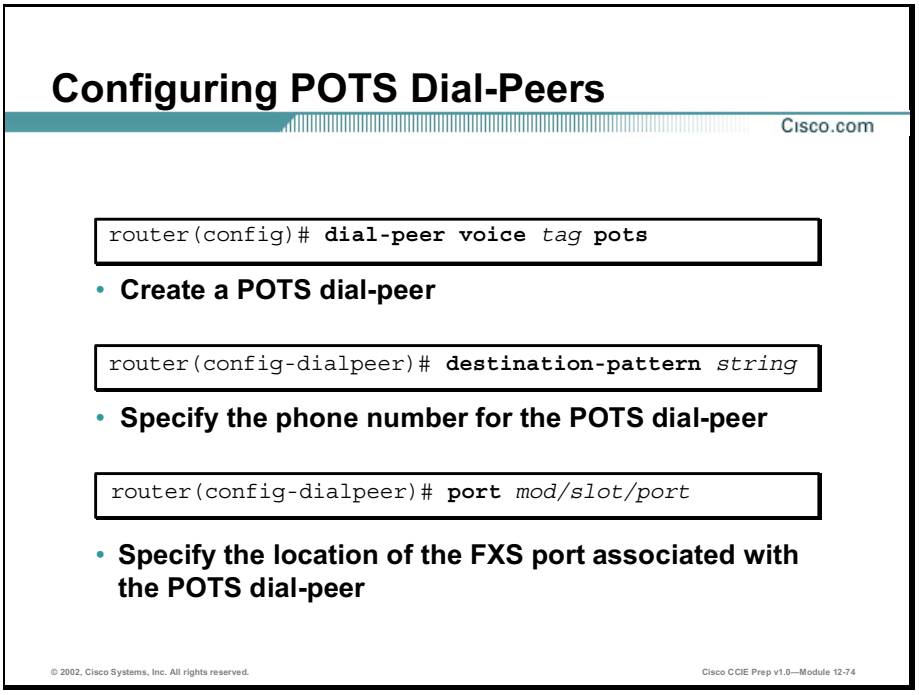

A Plain Old Telephone Service (POTS) dial-peer associates a phone number with a physical Foreign Exchange Station (FXS) port on a router. Before configuring a POTS dial-peer, determine the location of the router's FXS ports with the command: **show voice port summary**.

The output of this command will display the type of voice-ports (e.g., FXS, FXO, or E&M) installed on the router, in addition to the location of the FXS interface (e.g., mod/slot/port).

Once the FXS port and the associated phone number have been determined, the following steps detail how to set up a POTS dial-peer:

Step 1 From global configuration mode, enter the command **dial-peer voice** *tag* **pots**, where *tag* is a locally significant number unique among all dial-peers on the local router.

Step 2 From dial-peer configuration mode, enter the command **destination-pattern** *string*, where *string* is the phone number to be assigned to the phone.

Step 3 Still in dial-peer configuration mode, enter **port** *mod/slot/port*, where *mod/slot/port* specifies the module number, slot number, and port number of the FXS port. Note that on the MC3810, the port identifier only specifies a slot number and port number.

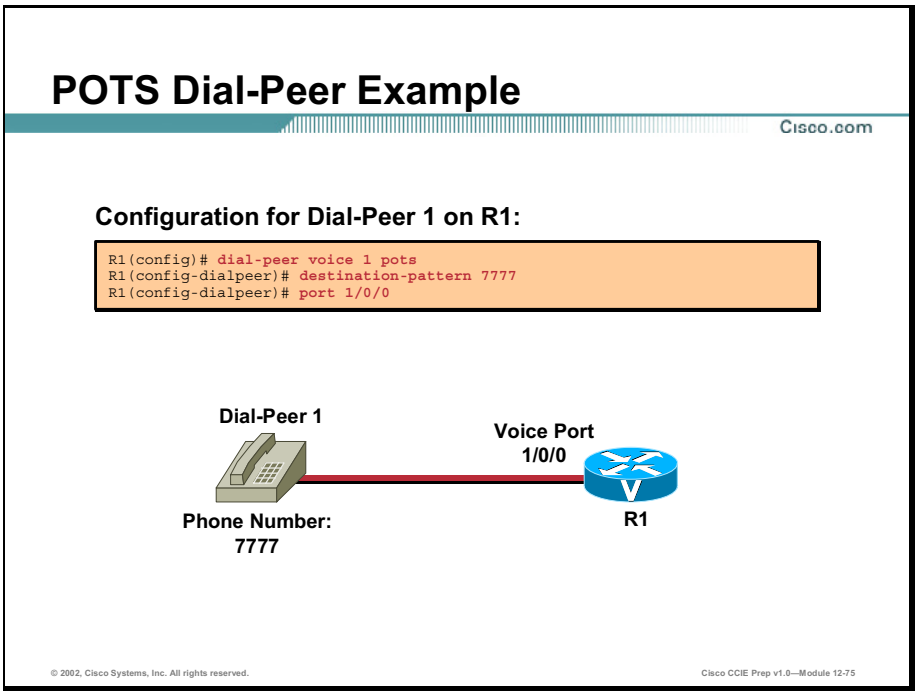

In the example shown, a phone with an extension number of 7777 is connected to an FXS port located in module 1, slot 0, port 0 of the router. The tag assigned to the dial-peer is 1.

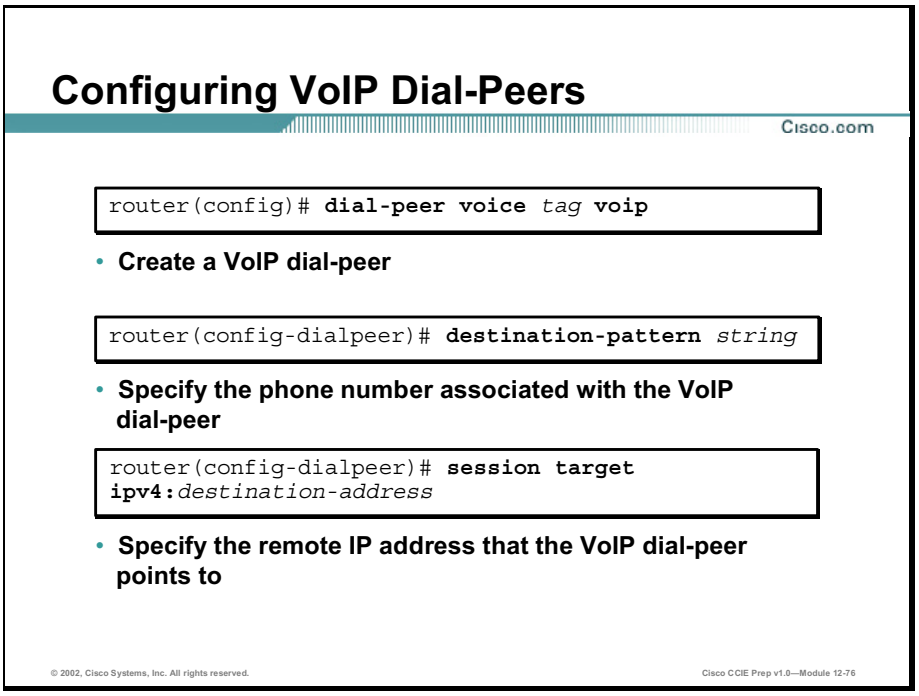

A VoIP dial-peer associates a phone number, or a pattern of phone numbers with the IP address of a remote router. The following steps detail how to configure a VoIP dial-peer:

Step 1 From global configuration mode, enter the command **dial-peer voice** ***tag*** **voip**, where ***tag*** is a locally significant number unique among all dial-peers on the local router.

Step 2 From dial-peer configuration mode, enter the command **destination-pattern** ***string***, where ***string*** is the phone number or pattern of phone numbers known to the remote router. For example, to specify all phone numbers in the range from 4000 – 4999, a pattern of **4…** could be specified. The **4…** pattern matches any four-digit phone number that begins with a 4.

Step 3 Still in dial-peer configuration mode, enter **session target ipv4:*destination-address***, where ***destination-address*** is the IP number of the remote router.

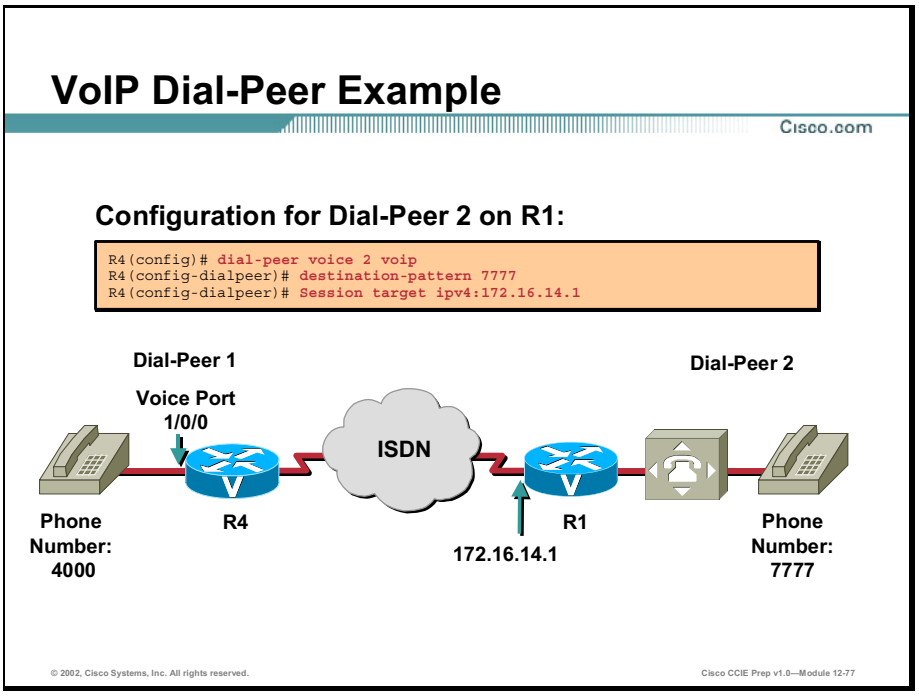

In the example shown, a phone with an extension number of 7777 is physically connected to R1 that has an IP address of 172.16.14.1. The VoIP dial-peer for R4 has a locally unique tag of 2.

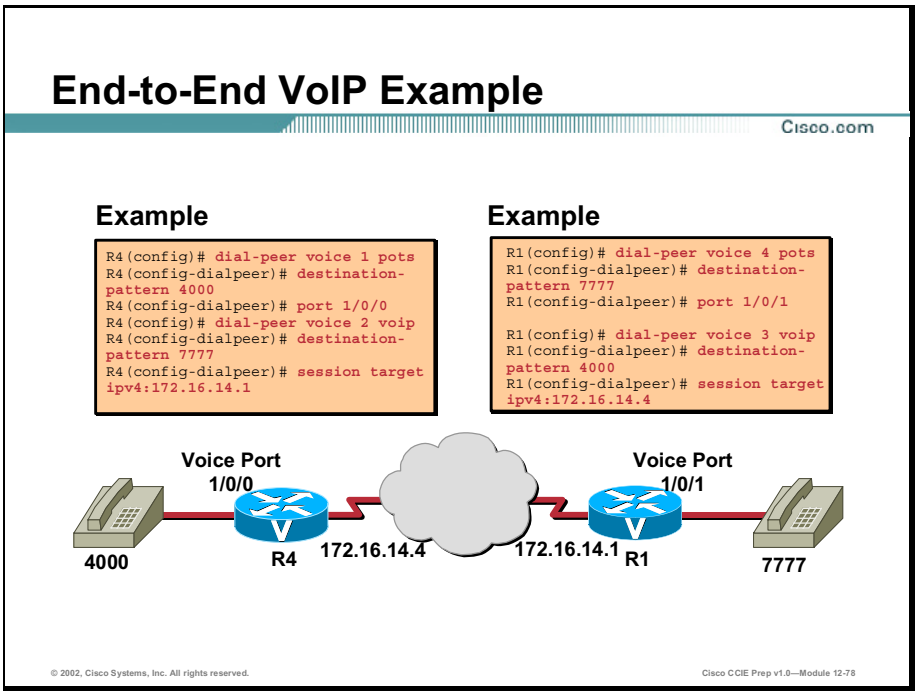

Consider an example, which combines POTS and VoIP dial-peers in a comprehensive end-to-end example. Here, extension 4000 can call extension 7777, and vice versa.

R4 Configuration:

```
R4(config)# dial-peer voice 1 pots
R4(config-dialpeer)# destination-pattern 4000
R4(config-dialpeer)# port 1/0/0
R4(config)# dial-peer voice 2 voip
R4(config-dialpeer)# destination-pattern 7777
R4(config-dialpeer)# session target ipv4:172.16.14.4
```

R1: Configuration

```
R1(config)# dial-peer voice 4 pots
R1(config-dialpeer)# destination-pattern 7777
R1(config-dialpeer)# port 1/0/1
R1(config)# dial-peer voice 3 voip
R1(config-dialpeer)# destination-pattern 4000
R1(config-dialpeer)# session target ipv4:172.16.14.4
```

Advanced VoIP Features

This section explains the Private Line Automatic Ringdown (PLAR) and Number Expansion (NUM-EXP) VoIP features.

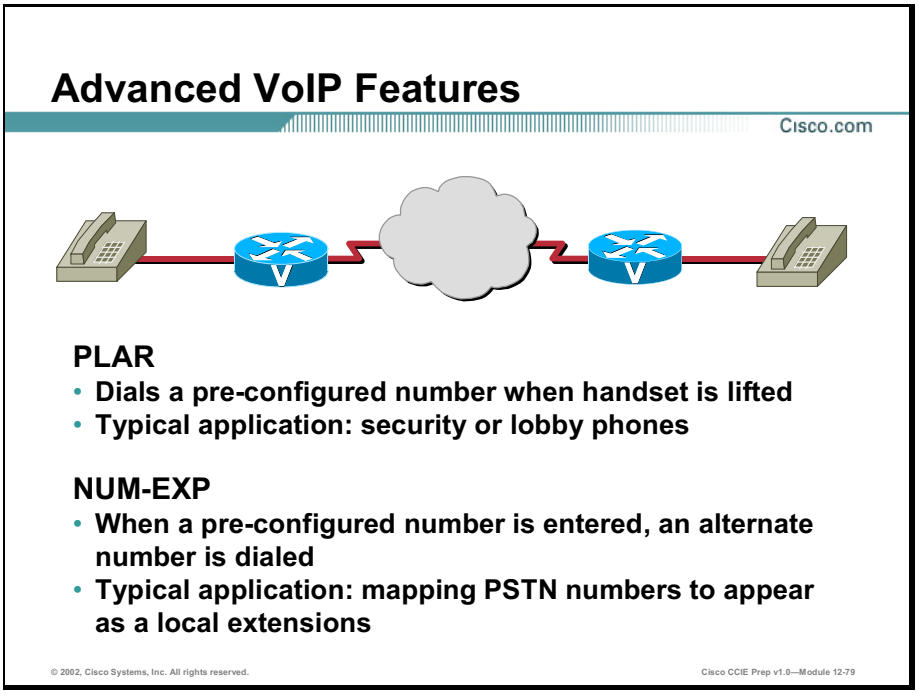

In addition to basic connectivity, Cisco routers offer a broad range of optional VoIP features. Two of these features, Private Line Automatic Ringdown (PLAR) and Number Expansion (NUM-EXP), will be highlighted in this section.

PLAR causes a predetermined extension number to be dialed immediately when someone picks up the handset on a phone. This is sometimes referred to as a “batphone”. Such an application may be useful for security phones in parking lots, for example. A caller could pick up the handset of the phone and immediately be connected to a representative in the Security department who could offer assistance.

NUM-EXP allows the Cisco router to translate between the number a caller enters and a specified number that is actually dialed. This feature can be useful when compiling a corporate directory. For example, a four-digit extension may be used to call internally within a corporation. If several employees work from home, the Number Expansion would be a good feature to use. This feature allows a four-digit extension to be published for these telecommuters and when an internal caller dialed one of those four-digit numbers, that number would be translated to a seven-digit number that used the Public Switched Telephone Network (PSTN). This number substitution would be transparent to the caller.

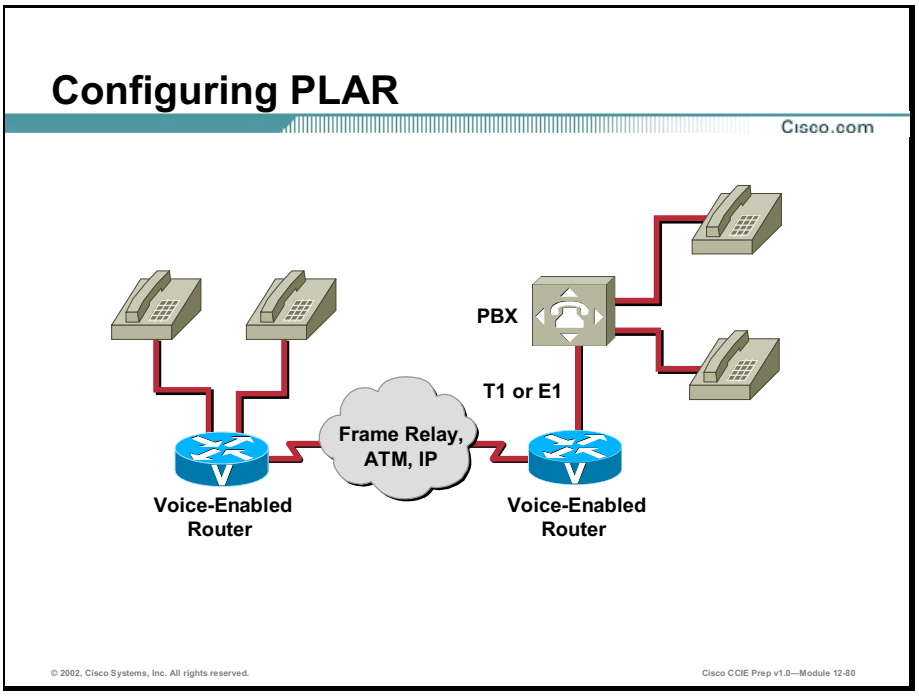

Before adding PLAR-specific commands to the router's configuration, the necessary POTS and VoIP dial-peers have to be set up. For example, if PLAR is directed to dial extension 4444 when a handset is lifted, then the router needs to have a corresponding dial-peer for extension 4444.

Once all of the dial-peers have been created, enter voice-port configuration mode, for the voice-port that will initiate the PLAR connection, with the global configuration command **voice-port** *mod/slot/port*, where *mod/slot/port* specifies the module number, slot number, and port number of the FXS port. Note that on the MC3810, the port identifier only specifies a slot number and port number.

From voice-port configuration mode, enter the command **connection plar** *number*, where *number* is the phone number to be dialed when the handset on this voice port is lifted. The voice port is now configured as a PLAR connection.

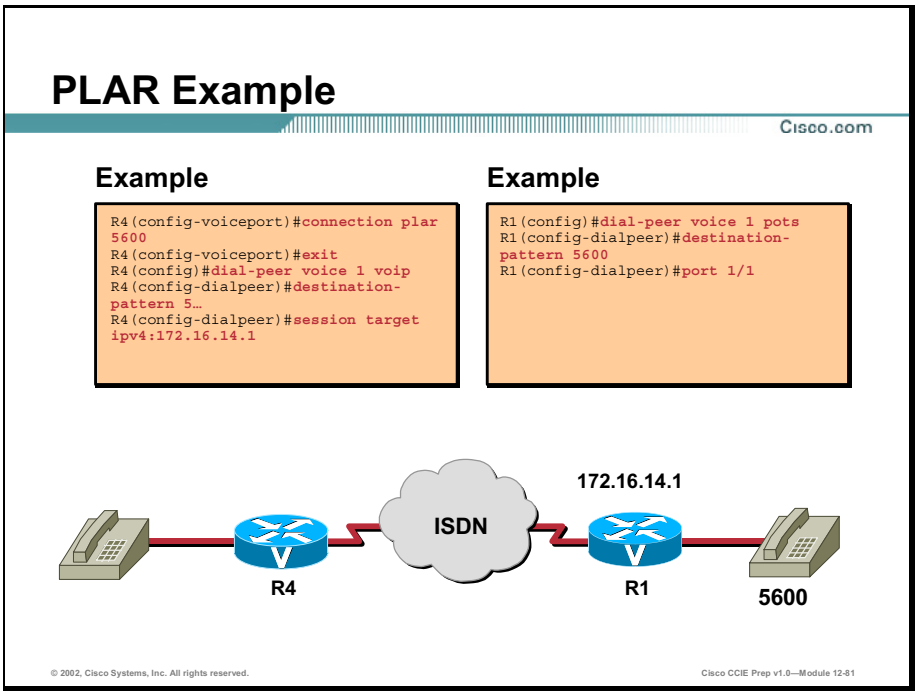

In this example, when the phone on the left has its handset lifted, it will immediately dial extension 5600. This is accomplished with the **connection plar 5600** command in voice-port configuration mode of R1.

R4 Configuration:

```
R4(config-voiceport)# connection plar 5600
R4(config-voiceport)# exit
R4(config)# dial-peer voice 1 voip
R4(config-dialpeer)# destination-pattern 5…
R4(config-dialpeer)# session target ipv4:172.16.14.1
```

R1 Configuration:

```
router(config)# dial-peer voice 1 pots
router(config-dialpeer)# destination-pattern 5600
router(config-dialpeer)# port 1/1
```

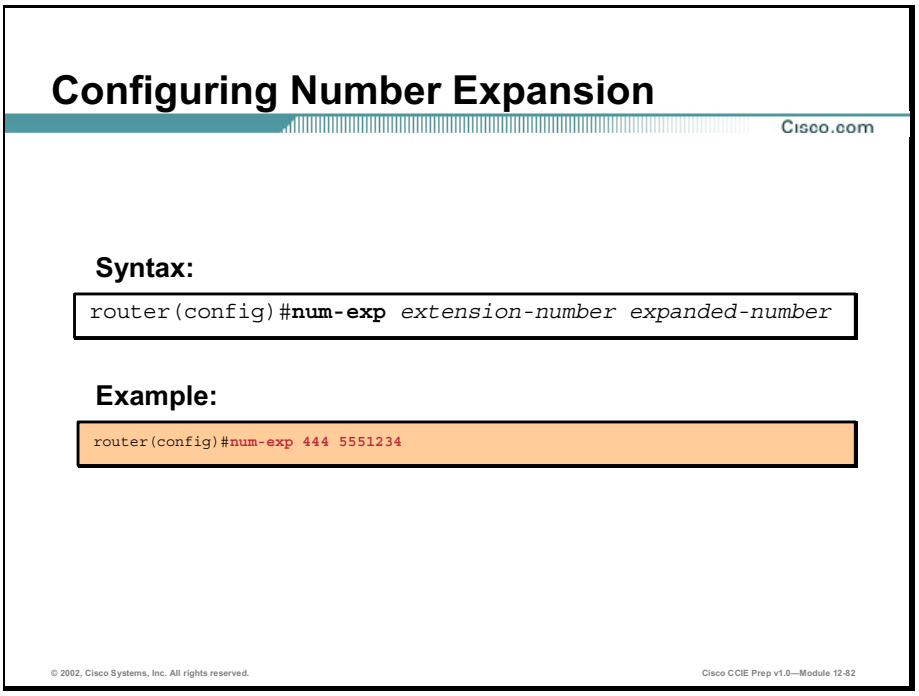

Number Expansion is configured from global configuration mode with the command **num-exp** *extension-number expanded-number*, where *extension-number* is the number entered by the caller, and *expanded-number* is the number that the router actually calls. In this example, when a caller dials 444, the router will initiate a call to 555-1234.

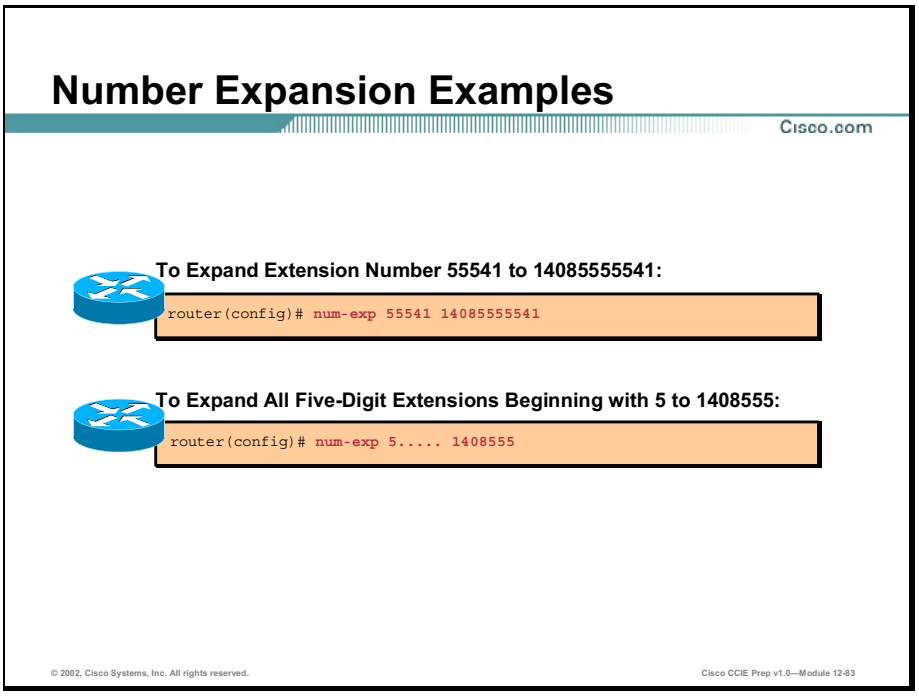

In the first example, if a caller dials the number 55541, the router will attempt to set up a call with 1-408-555-5541. Notice that hyphens are not entered in the syntax.

Example 1:

```
router(config)# num-exp 55541 14085555541
```

In the second example, if a caller dials any five-digit number beginning with the number 5, the router will attempt to set up a call with 1408555. Notice that the period wildcard is being used to specify any number in the range 50000-59999.

Example 2:

```
router(config)# num-exp 5..... 1408555
```

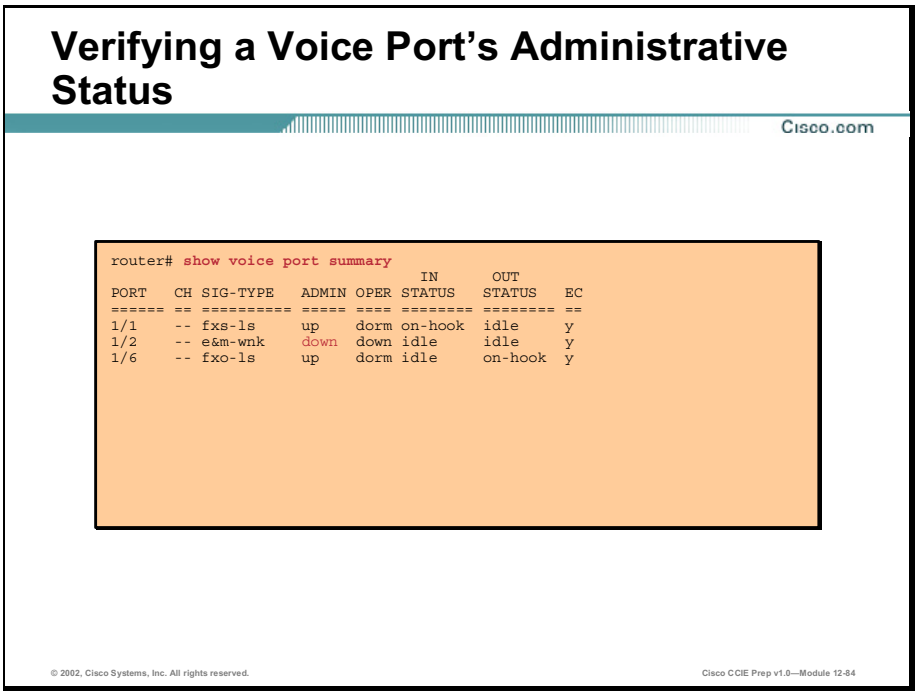

Like a network interface, a voice port can be in a shutdown state. To verify that the voice port is in an administratively up state, enter the command **show voice port summary**. In the example shown, notice that port 1/2 is administratively shutdown.

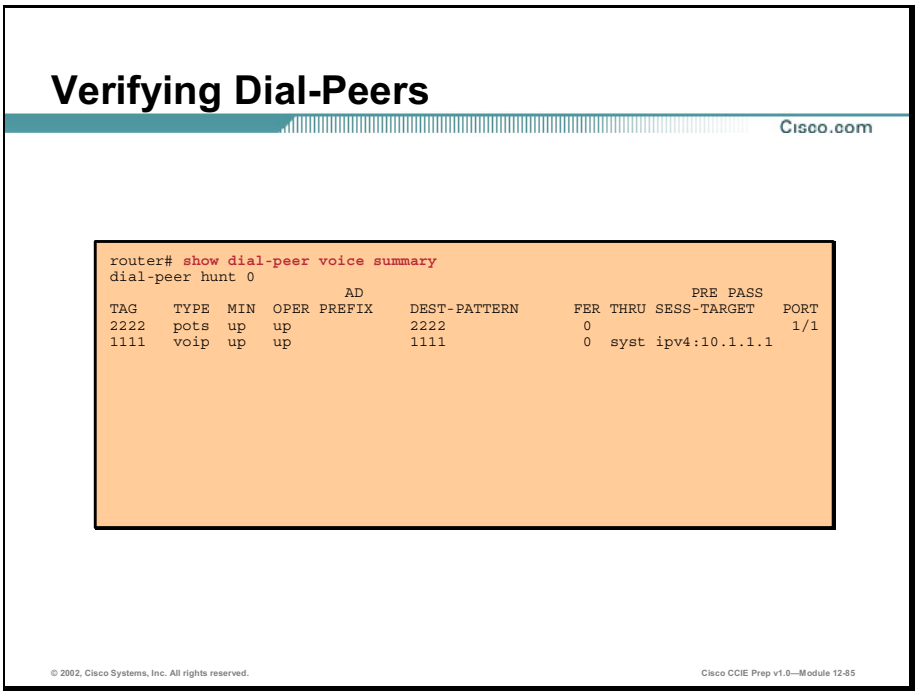

A summary of dial-peers can be viewed with the command **show dial-peer voice summary**. POTS dial-peer port assignment, VoIP dial-peer destination IP addresses, and associated phone numbers can be viewed from the results. Notice in the example that POTS dial-peer 2222 points to a physical FXS port of 1/1, and VoIP dial-peer 1111 points to a remote IP address of 10.1.1.1.

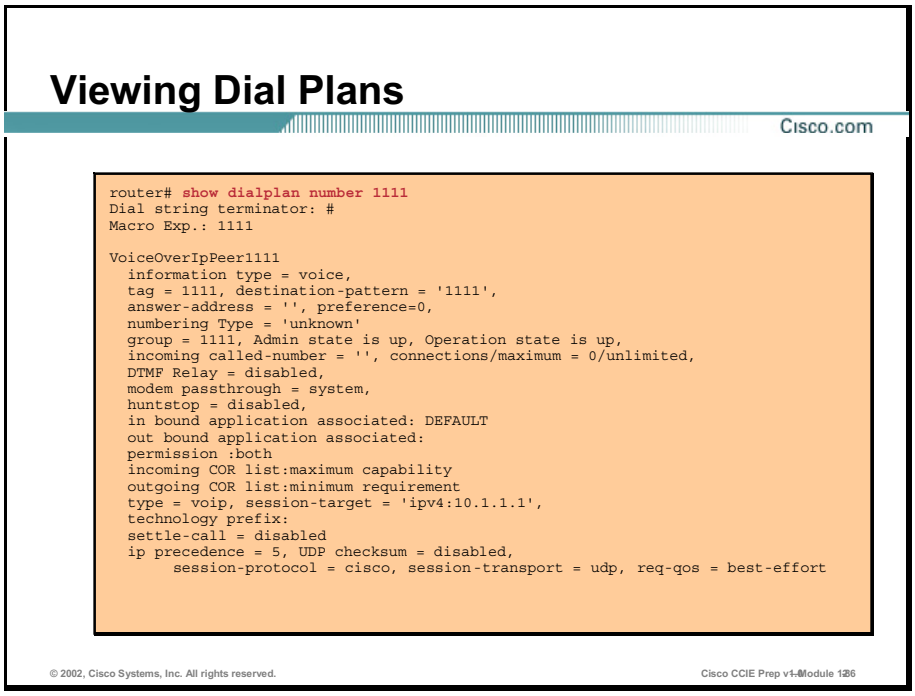

To verify that a router knows how to get to a specified destination phone number, use the command **show dialplan number** *ext_number*, where *ext_number* is the target phone number. In this example, the router does have information specifying how it can reach an extension number of 1111. A VoIP dial-peer pointing to a remote IP number of 10.1.1.1 is being used.

Summary

This section summarizes the key points discussed in this lesson.

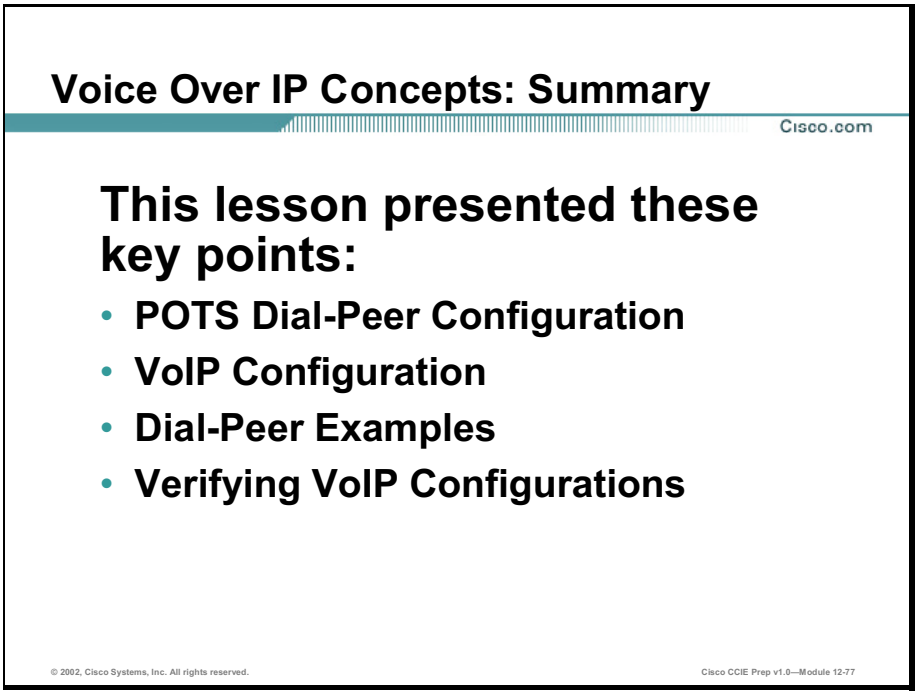

Next Steps

After completing this lesson, go to:

- Quality of Service Concepts

References

For additional information, refer to these resources:

- Cisco Voice over IP, Frame Relay, and ATM (CVOICE) Chapter 5
- VoIP Configuration Examples – http://www.cisco.com/univercd/cc/td/doc/product/access/acs_mod/1700/1750/1750voip/config.htm

Lesson Assessment (Quiz)

Q1) Which VoIP feature will cause a pre-configured phone number to be dialed when the handset of a phone is lifted?

A) POTS

B) PLAR

C) NUM-EXP

D) PSTN

Q2) Which dial-peer configuration command points to all numbers in the 4000 – 4999 range?

A) session target 4xxx

B) session target 4…

C) destination-pattern 4xxx

D) destination-pattern 4…

Q3) Which VoIP feature will cause a dialed number to be converted into a different number, which is used to actually place a call?

A) POTS

B) PLAR

C) NUM-EXP

D) PSTN

Q4) Which command configures a PLAR call to extension 5600?

A) router(config-voiceport)# **connection plar 5600**

B) router(config-dialpeer)# **connection plar 5600**

C) router(config-voiceport)# **plar 5600**

D) router(config-dialpeer)# **plar 5600**

Quality of Service Concepts

Overview

In a network environment with multiple types of data, some data may be more sensitive to delay than others. Cisco's Quality of Service (QoS) tools can influence which packets are forwarded first, and in the event of packet discard, some tools can influence which packets are discarded first. Additionally, policing and shaping tools can limit bandwidth usage.

Importance

QoS tools can be used to accomplish multiple objectives in the CCIE lab blueprint.

Objectives

Upon completing this lesson, you will be able to:

- Configure congestion management
- Configure Traffic Shaping
- Use FRTS as a shaping tool
- Use CAR as a policing tool
- Prevent congestion using WRED
- Verify QoS configuration

Learner Skills and Knowledge

To fully benefit from this lesson, you must have these prerequisite skills and knowledge:

- Passed the CCIE written qualification exam
- Completed the IP Quality of Service (IP QoS) course and/or the Deploying Quality of Service (DQoS) or have the equivalent knowledge

Outline

This lesson includes these sections:

- Overview
- Congestion Management
- Traffic Shaping
- Policing
- Congestion Avoidance
- QoS Verification
- Summary
- Lesson Assessment (Quiz)

Congestion Management

This section provides an overview of Cisco's queuing tools.

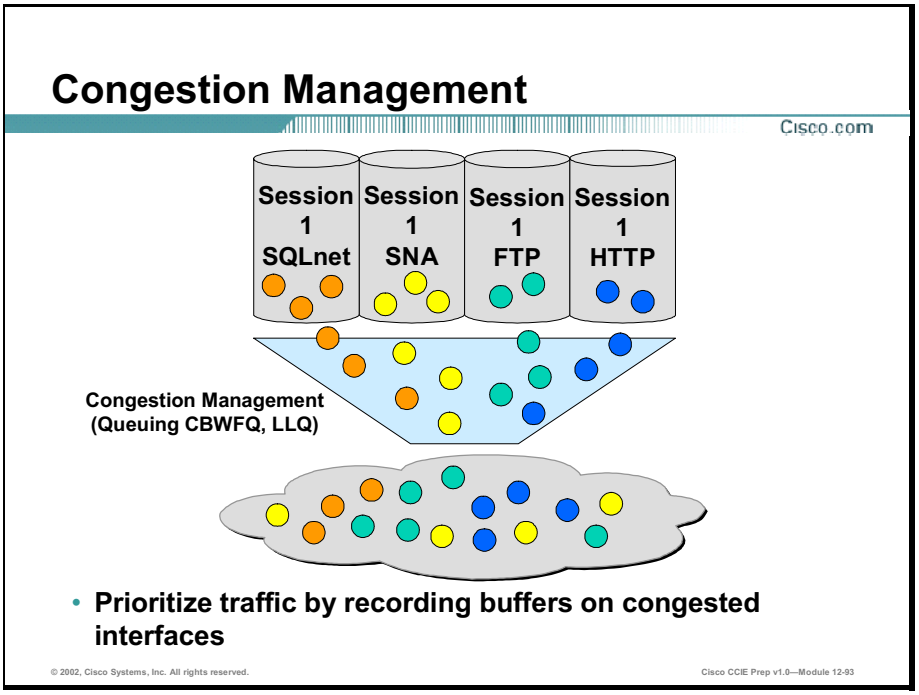

When multiple traffic flows are attempting to simultaneously exit an interface, Cisco typically uses First In First Out (FIFO) for high-speed interfaces, and Weighted Fair Queuing (WFQ) for interfaces running at less than 2 Megabits per second (Mbps).

Cisco's congestion management tools include several queuing methods, which can treat specified traffic with higher priority. This section will detail three of the more commonly used queuing tools: Priority Queuing (PQ), Custom Queuing (CQ), and Weighted Fair Queuing (WFQ).

Before Cisco's queuing tools can differentiate between traffic flows, the flows must first be classified. For example, a range of User Datagram Protocol (UDP) ports could be matched by an access list, and that access list could then be used to apply a specific queuing treatment to traffic matching that range of UDP ports.

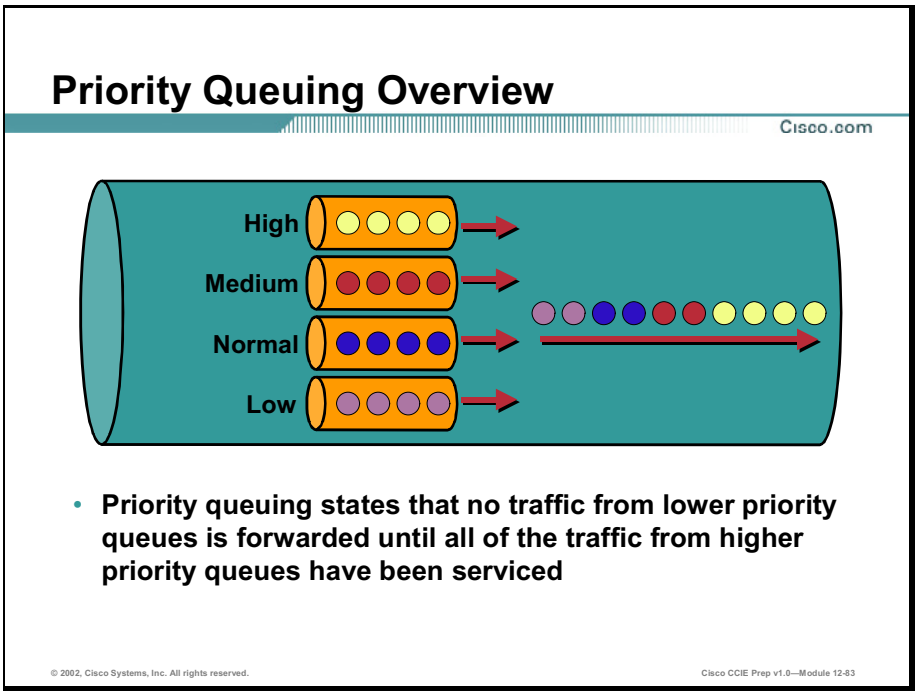

Priority Queuing (PQ) uses four predefined queues in which specified traffic types can be placed. The queues are then serviced in a "strict" fashion, in which no traffic from lower priority queues is forwarded until all of the traffic from higher priority queues has been serviced.

Specifically, PQ's four queues are defined as HIGH, MEDIUM, NORMAL, and LOW. Traffic in the medium queue must wait until all of the traffic in the high queue has been sent. Traffic in the normal queue has to wait until the high and medium queues have been completely emptied. Finally, only after the high, medium, and normal queues have been emptied is traffic from the low queue sent. While PQ can be useful in providing a strict level of high priority to traffic directed to the high queue, there is a potential for "starving out" flows that reside in lower priority queues.

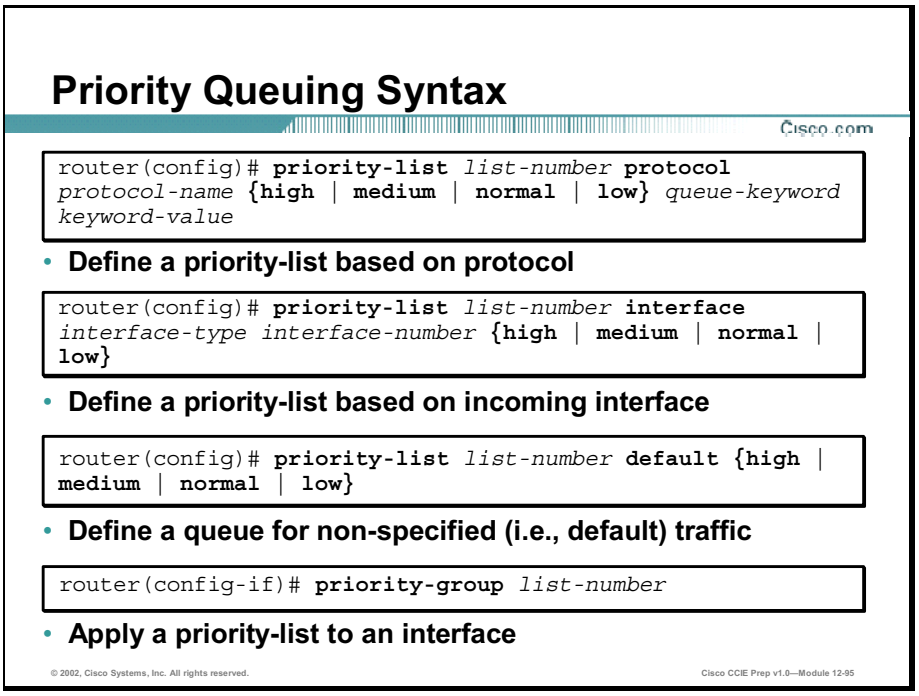

To define a priority list for a protocol, use the command **priority-list *list-number* protocol *protocol-name*** {**high** | **medium | normal** | **low**} ***queue-keyword keyword-value***, where ***list-number*** is a number from 1 – 16 that identifies the priority list, ***protocol-name*** specifies the Layer 2 or Layer 3 protocol (e.g., Internet Protocol (IP), AppleTalk, or Internetwork Packet Exchange (IPX)), and ***queue-keyword*** typical values include: **gt** for greater than, **lt** for less than, or **list** to specify an access list.

Alternately, traffic can be classified by the interface on which the traffic arrives, regardless of its protocol. To classify based on an incoming interface, use the command **priority-list *list-number* interface *interface-type interface-number*** {**high** | **medium | normal** | **low**}, where ***list-number*** is a number from 1 – 16 that identifies the priority list, ***interface-type*** specifies the name of the interface with incoming packets, and ***interface-number*** specifies the number (e.g., slot/port) of the incoming interface.

Traffic that is not specified by either of the other two methods is considered “default” traffic. To assign default traffic to a queue, use the command **priority-list *list-number* default** {**high** | **medium | normal** | **low**}, where ***list-number*** is a number from 1 – 16 that identifies the priority-list.

To apply the priority-list to an interface, use the command **priority-group *list***, where ***list*** is the number of the priority-list.

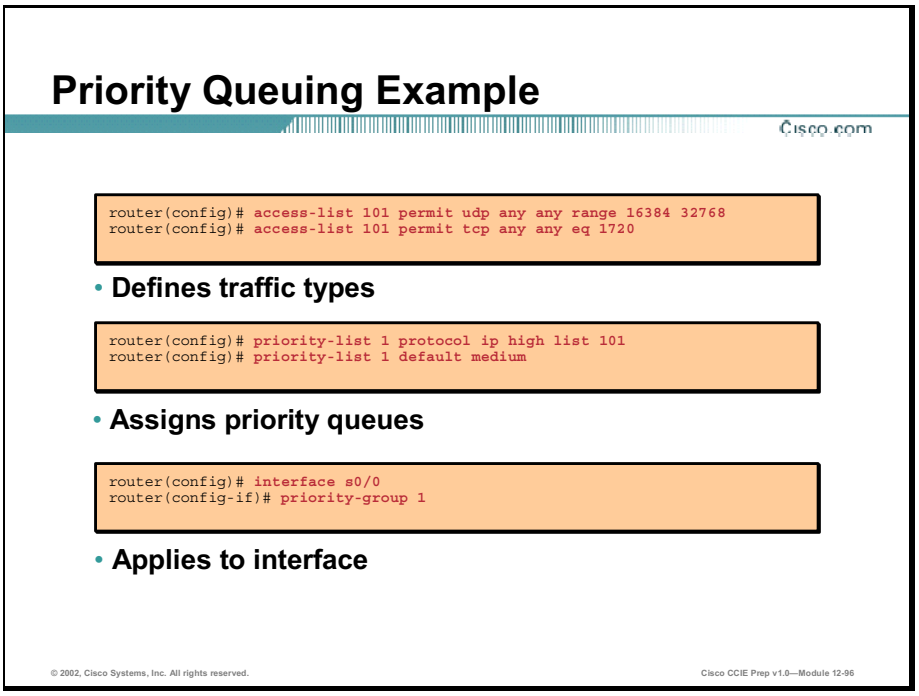

In the example shown, traffic with a destination port number in the UDP range of 16,384 – 32,768 and Transmission Control Protocol (TCP) traffic with a destination port number of 1720 is placed in the “high” queue. Remaining traffic (i.e., default traffic) is placed in the medium queue. The priority-list is then applied to the Serial 0/0 interface.

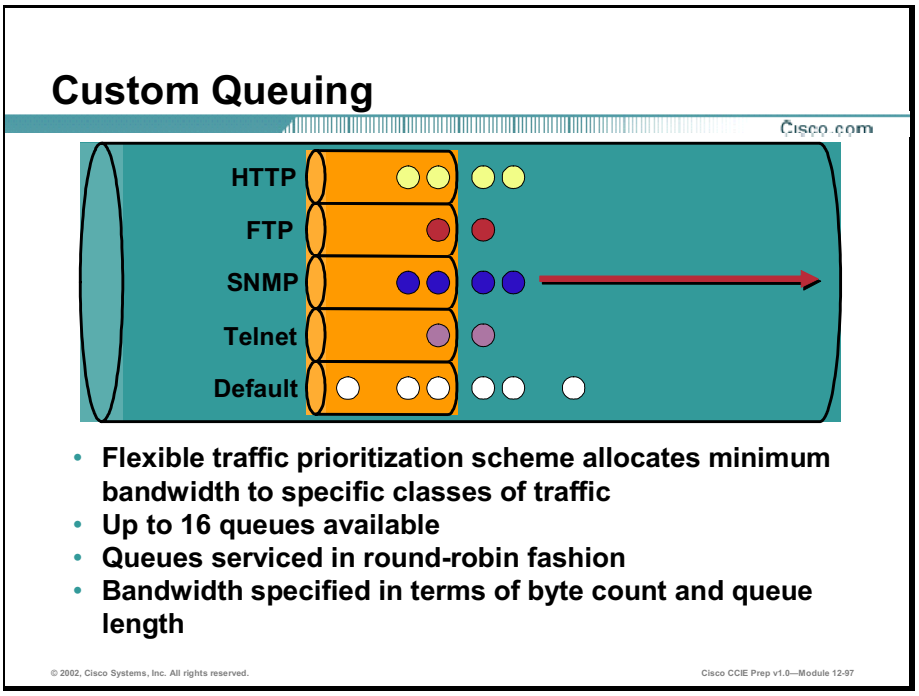

Custom Queuing (CQ) overcomes Priority Queuing's issue of protocol starvation by servicing a specified number of bytes from each of the defined queues. With CQ, traffic is classified and separated in up to 16 separate queues. The CQ scheduler then removes a certain number of bytes from each queue, as it services all of the queues in a round-robin fashion.

By specifying the number of bytes to be removed from each queue during each round-robin cycle, CQ can be configured to allocate a percentage of bandwidth to each type of traffic. For example, to make 30% of the bandwidth available for File Transfer Protocol (FTP) traffic, 20% of the bandwidth available for Telnet traffic, 25% of the bandwidth available for Hypertext Transport Protocol (HTTP) traffic, and 25% of the bandwidth available for default (i.e., unspecified) traffic, the number of bytes to be serviced from each queue could be divided as follows:

- FTP – 3000 Bytes
- Telnet – 2000 Bytes
- HTTP – 2500 Bytes
- Default – 2500 Bytes

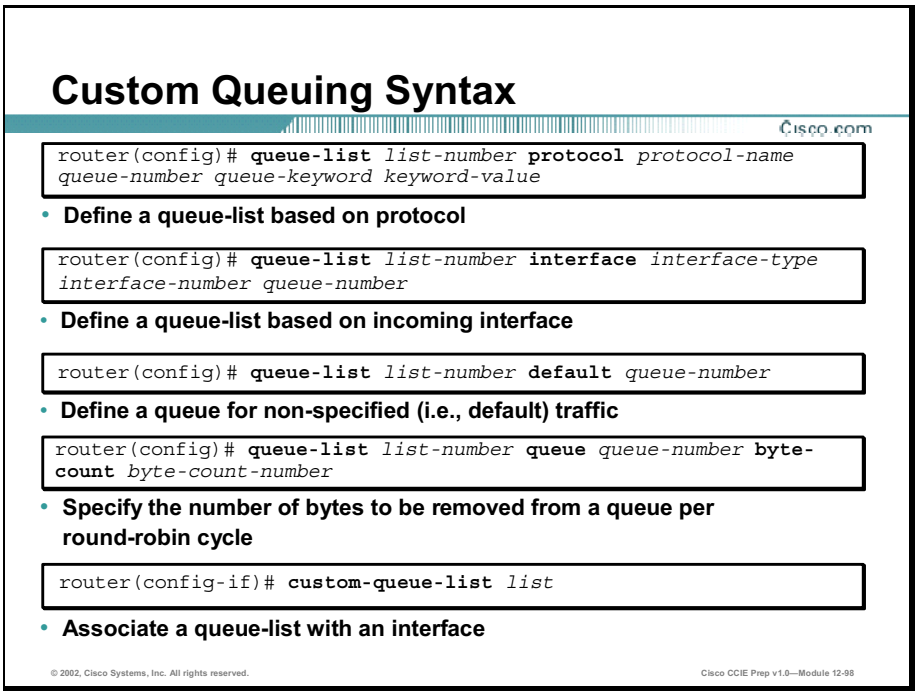

To include traffic from a particular protocol or traffic defined by an access list in a queue, use the command **queue-list** ***list-number*** **protocol** ***protocol-name queue-number queue-keyword keyword-value***, where ***list-number*** is the number of the queue list in the range of 1 – 16, ***protocol-name*** is the protocol type (e.g., IP, AppleTalk, or IPX), ***queue-number*** is the number of the queue in the range of 1 – 16, and where typical ***queue-keywords*** include **gt** for greater than, **lt** for less than, or **list** to specify an access list.

To classify traffic by the incoming interface, use the command **queue-list** ***list-number*** **interface** ***interface-type interface-number queue-number***, where ***list-number*** is the number of the queue list in the range of 1 – 16, ***interface-type*** specifies the name of the interface with incoming packets, ***interface-number*** specifies the number (e.g., slot/port) of the incoming interface, and ***queue-number*** specifies the number of the queue in the range of 1 – 16.

Traffic that is not specified by either of the other two methods is considered "default" traffic. To assign default traffic to a queue, use the command **queue-list** ***list-number*** **default** ***queue-number***, where ***list-number*** is the number of the queue list in the range of 1 – 16, and ***queue-number*** specifies the number of the queue in the range of 1 – 16.

To assign a specified number of bytes to a queue, use the command **queue-list** ***list-number*** **queue** ***queue-number*** **byte-count** ***byte-count-number***, where ***list-number*** is the number of the queue list in the range of 1 – 16, ***queue-number*** specifies the number of the queue in the range of 1 – 16, and ***byte-count-number*** specifies the minimum number of bytes to be removed from a queue per round-robin cycle. The default byte-count is 1,500 bytes.

To apply the queue list to an interface, use the command **custom-queue-list** ***list***, where ***list*** is the number of the queue list.

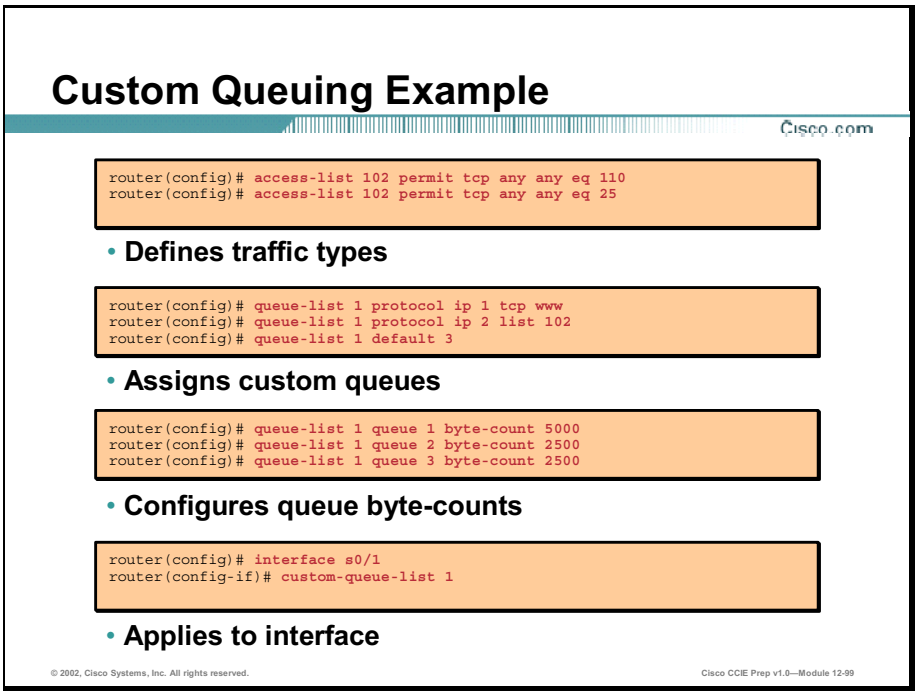

In the example shown, 25% of the bandwidth is being reserved for e-mail (i.e., SMTP and POP3) traffic; 50% of the bandwidth is being reserved for web traffic; and 25% of the bandwidth is being reserved for default (i.e., unspecified) traffic. This is based on a total available bandwidth of 10,000 bytes. The custom queuing configuration is then applied to the Serial 0/0 interface.

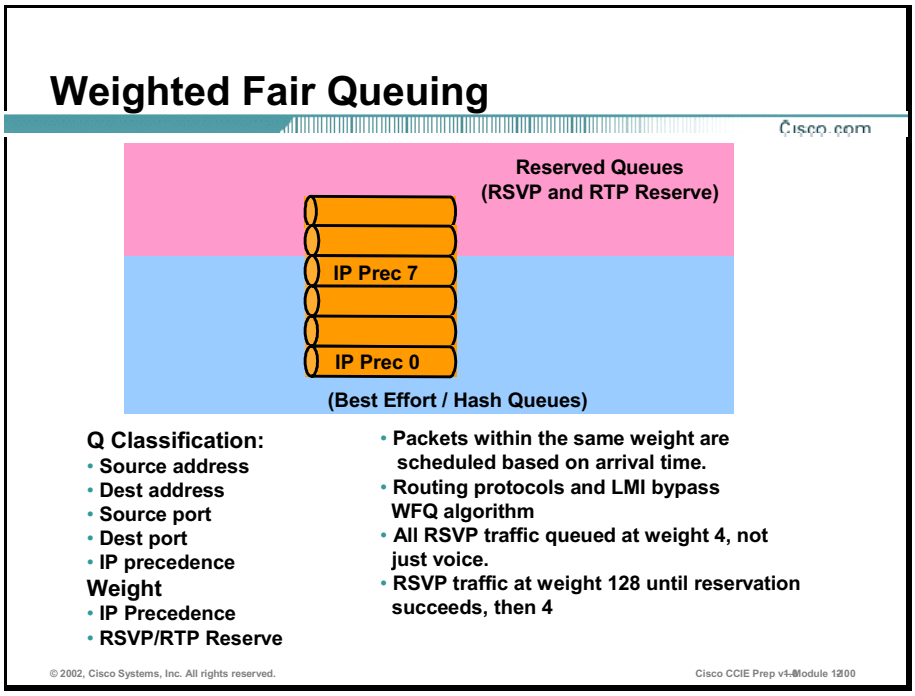

Weighted Fair Queuing (WFQ) attempts to share the bandwidth fairly between low-bandwidth and high-bandwidth flows. For example, if there were simultaneous FTP and Telnet sessions, WFQ would not allow the low-bandwidth Telnet session to be overwhelmed by the high-bandwidth FTP session.

If a 64-byte packet and a 256-byte packet arrived at the queue at the same time, the WFQ scheduler would send the 64-byte packet first. However, assigning a "weight" to a particular flow can alter this "fair queuing" behavior. WFQ considers the IP Precedence value of a flow when determining each packet's sequence number. Specifically, the weight of a packet is 4,096/(IP Prec. +1) or 32,768/(IP Prec. + 1), depending on Internetwork Operating System (IOS) version. This weight is used along with the packet size to determine the order in which the WFQ scheduler will service packets.

When used in conjunction with Reservation Protocol (RSVP), WFQ gives traffic in an RSVP flow extremely high priority. Instead of having a weight of 32,768/(IP Prec. + 1), an RSVP packet has a weight of 4 or 6, depending on the IOS version. Therefore, even if the RSVP packets are large, they will be scheduled with a greater preference than non-RSVP flows.

WFQ is Cisco's default queuing method for low-speed interfaces (i.e., less than 2 Mbps). Since WFQ is a default setting, the command **fair-queue** does not appear when viewing the running configuration.

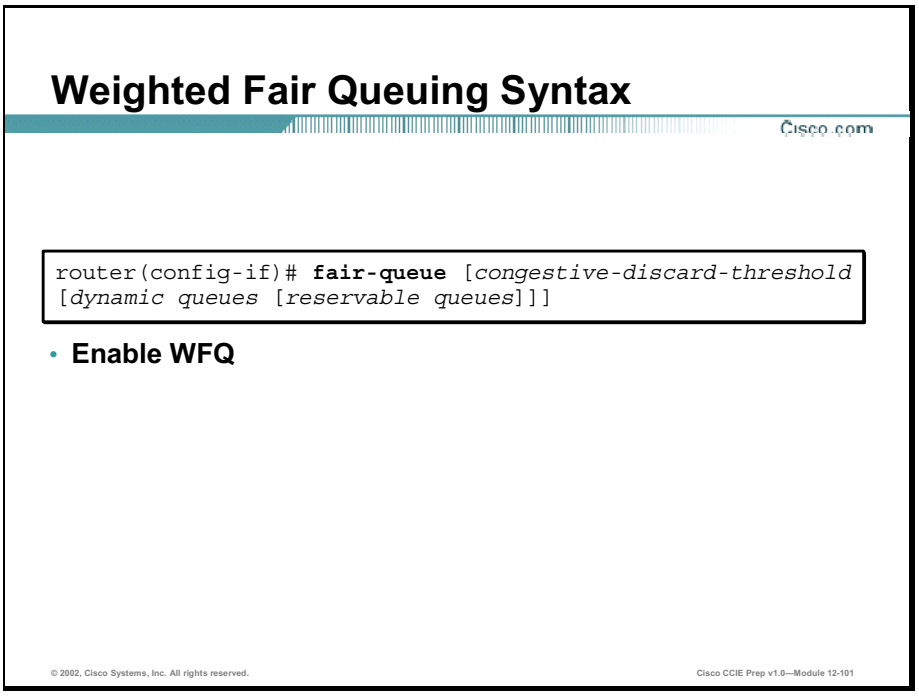

To enable WFQ on an interface, use the command **fair-queue** [***congestive-discard-threshold*** [***dynamic queues*** [***reservable queues***]]], where ***congestive-discard-threshold*** is the number of messages allowed in each queue, ***dynamic queues*** is the number of dynamic queues to be used for best-effort conversations, and ***reservable queues*** is the number of reservable queues used for reserved conversations (e.g., RSVP) in the range 0 – 1000.

The ***congestive-discard-threshold*** parameter applies only to high-volume conversations that have more than one message in the queue. If an individual conversation queue contains more messages than the congestive discard threshold, that conversation will not have any new messages queued until that queue's content drops below one-fourth of the congestive discard value.

Traffic Shaping

This section details the procedures for configuring Frame Relay Traffic Shaping (FRTS).

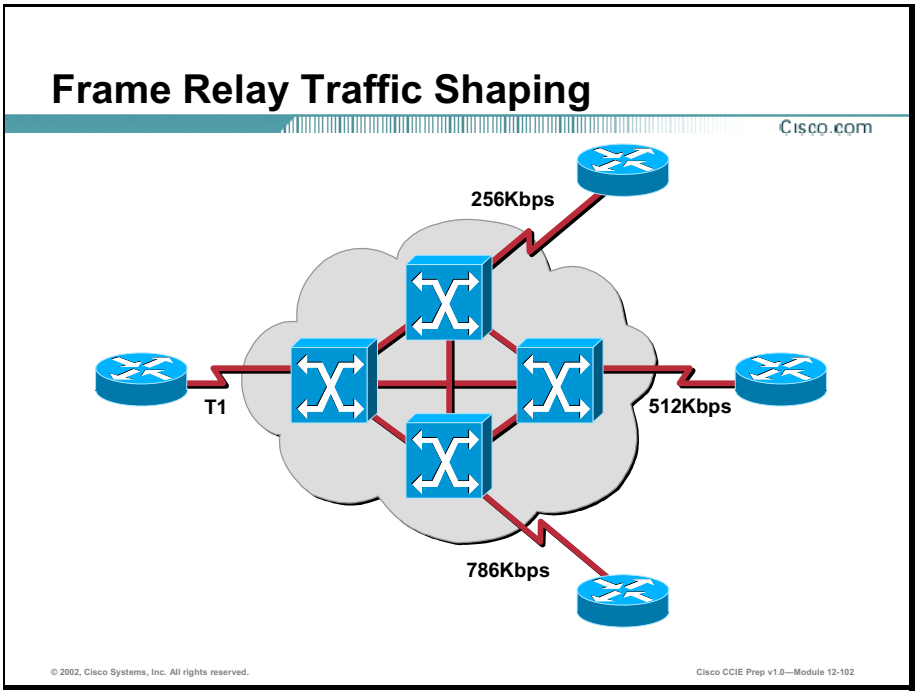

Since Frame Relay uses Permanent Virtual Circuits (PVCs), it is possible for one end of the PVC to have a faster access rate than the other end of the PVC. This condition can lead to oversubscription. Frame Relay Traffic Shaping (FRTS) can limit the bandwidth on PVCs to prevent oversubscription.

The parameters used by FRTS include:

- **Committed Information Rate (CIR):** The CIR is the average amount of traffic the service provider has guaranteed to forward over a specified amount of time.

- **Committed Burst (Bc):** The Bc is the amount of traffic that the service provider will attempt to send during any time interval. However, if there is congestion on the network, traffic above the Bc range will be discarded first, because those frames have their Discard Eligibility (DE) bit set to a 1.

- **Excess Burst (Be):** The Be is the amount of traffic above the Bc that the service provider will attempt to send during a given time interval. If a router could have sent traffic up to the Bc level during a previous time interval, but it did not have traffic to send at that moment, then that unused bandwidth can be “banked” and used in a future time interval, up to the Be limit.

- **Committed Time Interval (Tc):** The Tc is a measurement of time in which the service provider will attempt to burst an amount of data equal to Bc.

- **Backward Explicit Congestion Notification (BECN):** A BECN notifies the sending router that congestion is being experienced in the cloud and requests that the router slow down its transmission, below its CIR.

- **Minimum CIR (MINCIR):** The MINCIR is the minimum value to which the CIR will drop in the presence of BECN notifications from the service provider.

The mathematical relationship between the CIR, Bc, and Tc is critical to manipulating the duration of the timing interval. This relationship can be stated as: Tc=Bc/CIR.

If Tc is not specified, then the maximum value of Tc is 1/8 of a second or 125 ms. By knowing the CIR and the maximum value of Tc, then Bc can be calculated.

The relationship between Bc and Be can be further defined as Bc+Be=PVC based on the time interval or Tc.

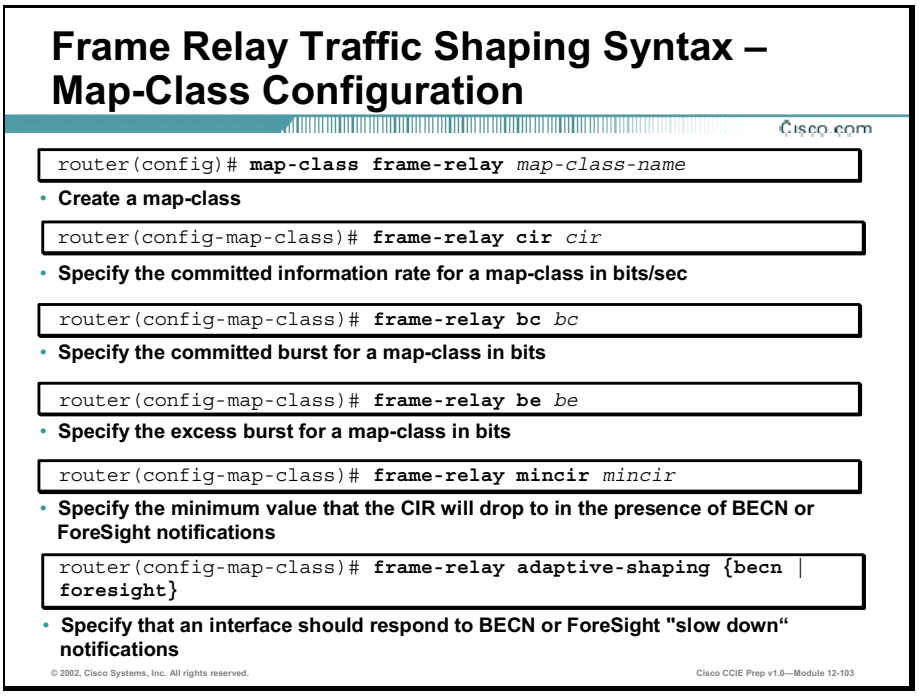

FRTS parameters are configured within a map-class. The map-class is then applied to one of the frame-relay Data Link Connection Identifiers (DLCIs).

To create a map-class, use the command **map-class frame-relay** ***map-class-name***, where ***map-class-name*** is a locally unique name that will be used to associate the map-class with a frame relay DLCI. This command will enter the router into map-class configuration mode.

To specify the CIR, in bits per second, enter the command **frame-relay cir** ***cir***, where ***cir*** is the agreed upon average number of bits to be sent in a second.

To specify the committed burst, in bits, enter the command **frame-relay bc** ***bc***, where ***bc*** is the number of bits above the CIR that the service provider will attempt to send during the Tc timing interval.

To specify the excess burst, in bits, enter the command **frame-relay be** ***be***, where ***be*** is the number of bits that the service provider will attempt to send above the bc level, when the router has not fully used the bc level during the previous time intervals.

To specify the minimum value that the CIR will drop to in the presence of BECN or ForeSight "slow down" notifications, use the command **frame-relay mincir** ***mincir***, where ***mincir*** is the minimum traffic rate in bits per second.

To specify that the DLCI associated with this map-class should respond to BECN or ForeSight "slow down" notifications, use the command **frame-relay adaptive-shaping** {**becn** | **foresight**}.

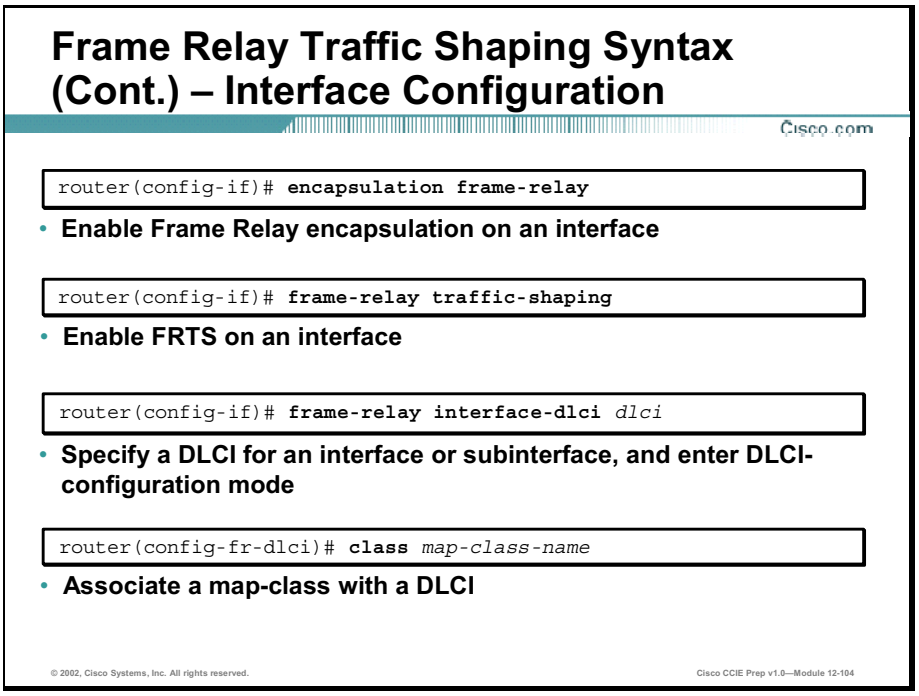

The first step in configuring FRTS on an interface is to configure the interface for Frame Relay encapsulation, with the command **encapsulation frame-relay**.

To enable traffic shaping on a frame relay interface, use the command **frame-relay traffic-shaping**.

The previously configured map-class specifies FRTS parameters. To associate a map-class with a DLCI, enter DLCI-configuration mode with the command **frame-relay interface-dlci *dlci***, where ***dlci*** is the locally significant PVC identifier.

Once in DLCI-configuration mode, the map-class can be associated with the DLCI using the command **class *map-class-name***, where ***map-class-name*** is the name of a map-class configured from global configuration mode.

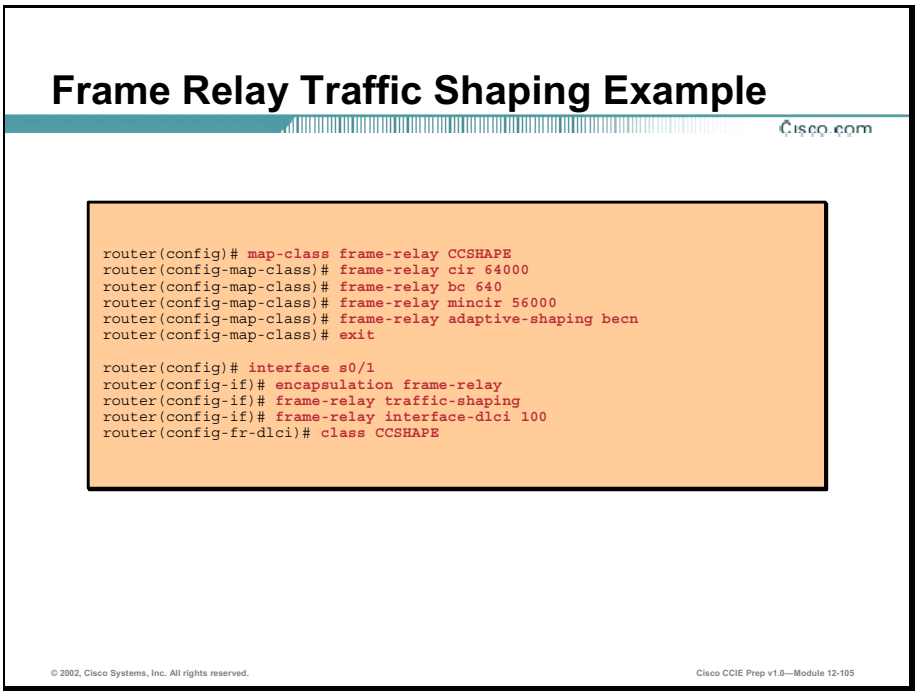

In the example shown, interface Serial 0/1 is configured for FRTS. The map-class CCSHAPE is used to specify the following characteristics for DLCI 100:

- The CIR is set to 64 kilobits per second (kbps)
- Bc is set to a value that causes the Tc interval to be 10 ms (Tc=Bc/CIR)
- The DLCI should respond to BECN messages
- In the presence of BECNs, the CIR should not drop below 56 kbps

Policing

This section details the procedures for configuring Cisco's Committed Access Rate (CAR) policing feature.

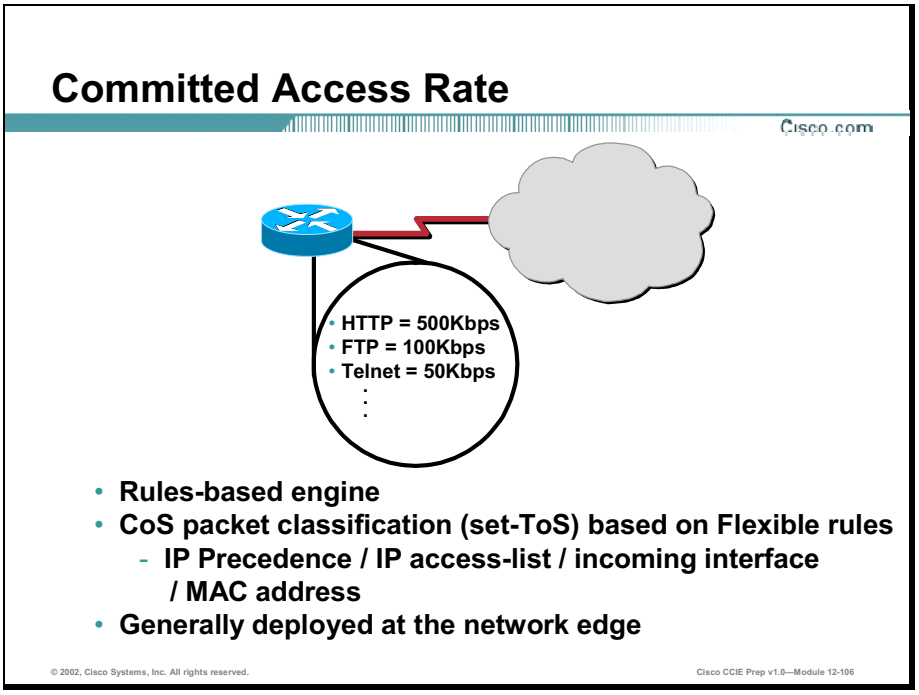

Committed Access Rate (CAR) is a policing tool that limits the amount of bandwidth given to a particular traffic classification. Additionally, CAR can mark traffic with IP Precedence values.

Typically, traffic is classified using an access list, and if the bandwidth of the traffic matching the access list is less than the specified CIR, then a "conform" action is performed. If the bandwidth of the traffic matching the access list is greater than the CIR, an "exceed" action is performed. CAR also supports the policing of an interface, not just traffic that matches an access list.

The following are CAR's supported conform and exceed actions:

- **Transmit -** Send the packet
- **Drop -** Discard the packet
- **Continue -** Go to the next CAR rate-limit statement in the list
- **Set Precedence and Transmit -** Rewrite the IP Precedence value in the packet's Type of Service (ToS) byte to the specified value, and send the packet
- **Set Precedence and Continue -** Rewrite the IP Precedence value in the packet's ToS byte to the specified value, and go to the next CAR rate-limit statement in the list

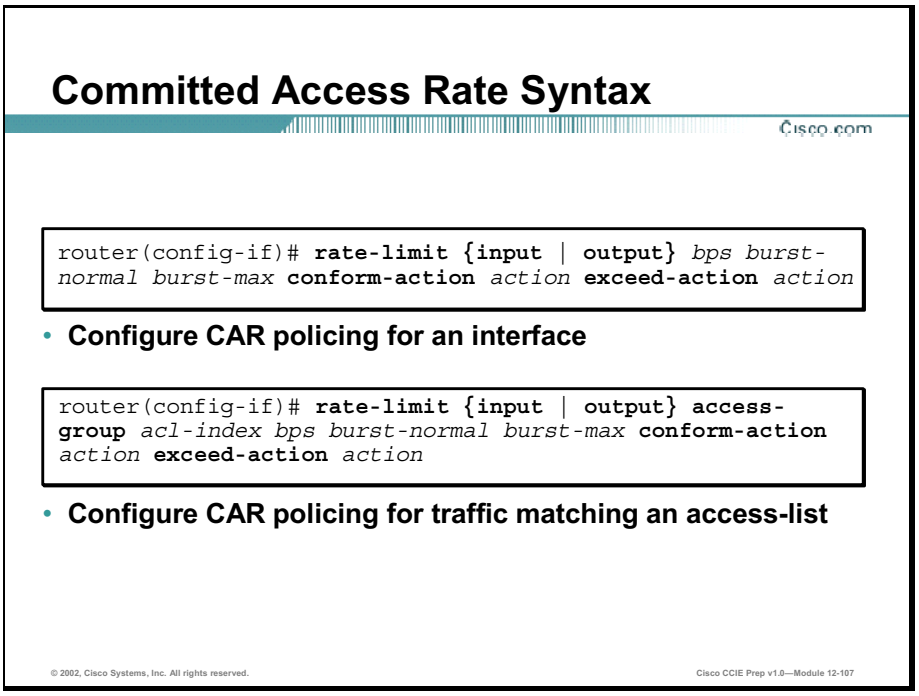

To configure rate limiting on an entire interface, use the command **rate-limit** {**input** | **output**} ***bps burst-normal burst-max*** **conform-action** ***action*** **exceed-action** ***action***, where ***bps*** is the CIR to which CAR is policing in bits per second, ***burst-normal*** is the committed burst per time interval in bytes, ***burst-max*** is the excess burst per time interval in bytes, and ***action*** can be one of the following:

- **continue:** Evaluate the next **rate-limit** command
- **drop:** Drop the packet
- **set-prec-continue** ***new-prec***: Set the IP Precedence, and continue to the next **rate-limit** command
- **set-prec-transmit** ***new-prec***: Set the IP Precedence, and send the packet
- **transmit**: Send the packet

To configure rate limiting for traffic specified by an access list, use the command **rate-limit** {**input** | **output**} **access-group** ***acl-index bps burst-normal burst-max*** **conform-action** ***action*** **exceed-action** ***action***, where ***acl-index*** is the access list number that identifies the traffic to be policed.

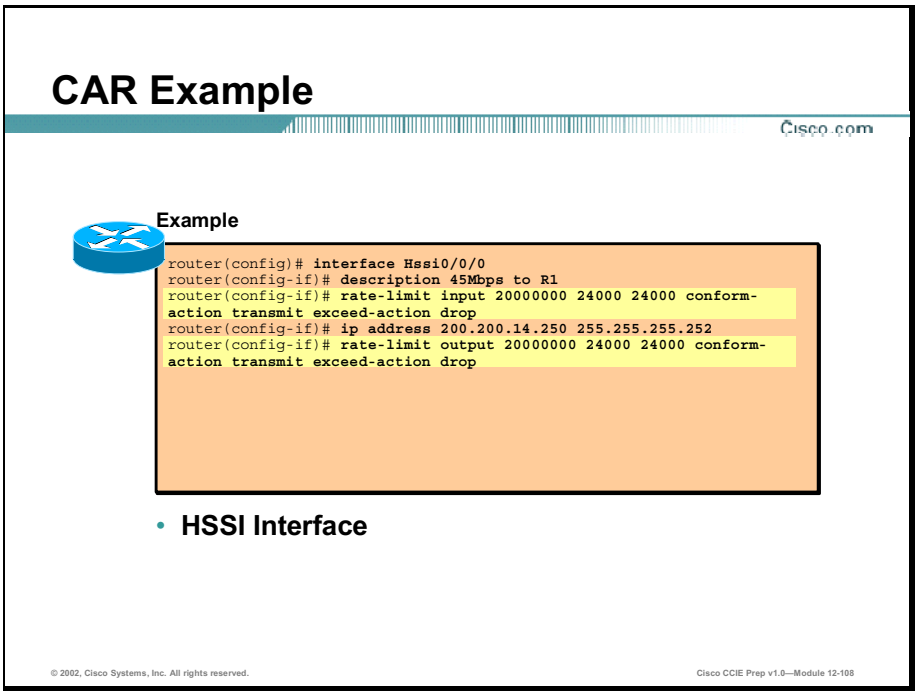

In the example shown, the High-Speed Serial Interface (HSSI) is being policed to 20 Mbps, with a burst capability of 24 kilobytes per time interval. If traffic on the interface conforms to the 20 Mbps rate limit, then the traffic is passed, due to the **transmit** conform action. If traffic exceeds the 20 Mbps rate limit, then the traffic is dropped, due to the **drop** exceed action.

Also, notice that there are two **rate-limit** commands in this example. One is applied inbound on the interface, while one is applied outbound on the interface. While traffic shaping is used on outbound traffic only, policing can be used inbound, outbound, or both simultaneously.

Congestion Avoidance

This section explains the configuration of Weighted Random Early Detection (WRED).

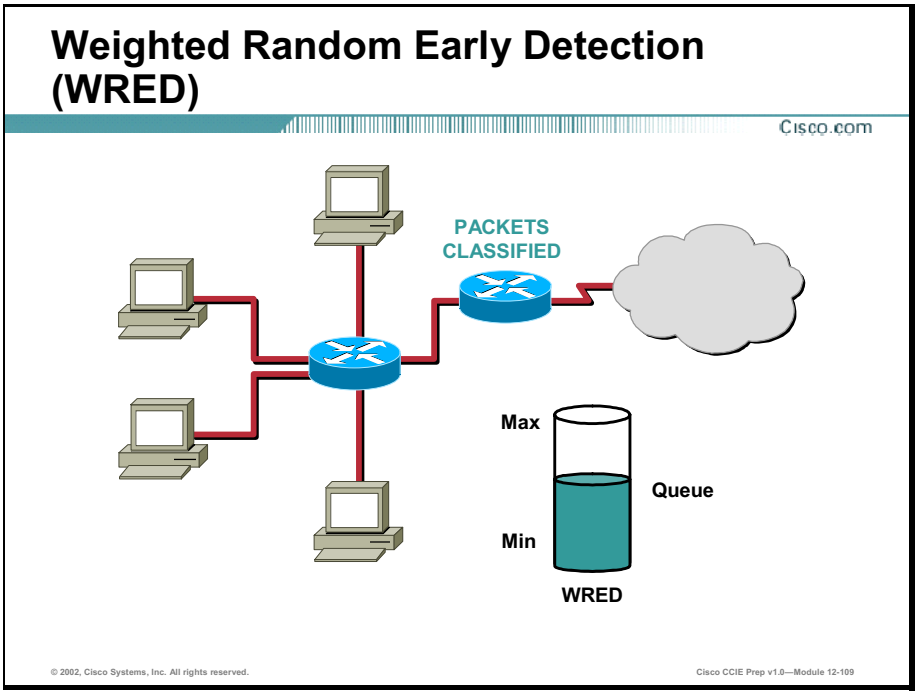

While queuing tools provide congestion management, there is still a need for congestion avoidance, especially for TCP flows. When congestion occurs on an output queue, the default behavior is to "tail drop" packets. Tail dropping packets involves the discarding of packets attempting to enter the queue after the queue is full.

The result of a queue filling up is the simultaneous tail drop behavior of all flows attempting to enter the queue. This action causes all of the TCP flows to enter TCP Slow Start, which is the reduction of the TCP window size to 1. The window size then increases exponentially up to half of the original congestion window size. At that point, the window size increases linearly. This phenomenon of multiple TCP flows simultaneously entering TCP Slow Start is called Global Synchronization, which results in unused bandwidth, due to the small window size of all TCP flows.

Therefore, a mechanism is needed to prevent Global Synchronization. An industry-standard method, called Random Early Detection (RED), will begin to randomly discard packets as the queue nears capacity. Since the queue never fills to capacity, due to the random discarding of packets, Global Synchronization never occurs.

One negative aspect of RED, however, is that it does not distinguish between flows. Cisco's implementation takes RED a step further, and discards packets at a rate depending on the IP Precedence or Differentiated Services Code (Control) Point (DSCP) values. Support for DSCP-based WRED was introduced in IOS 12.1(5)T.

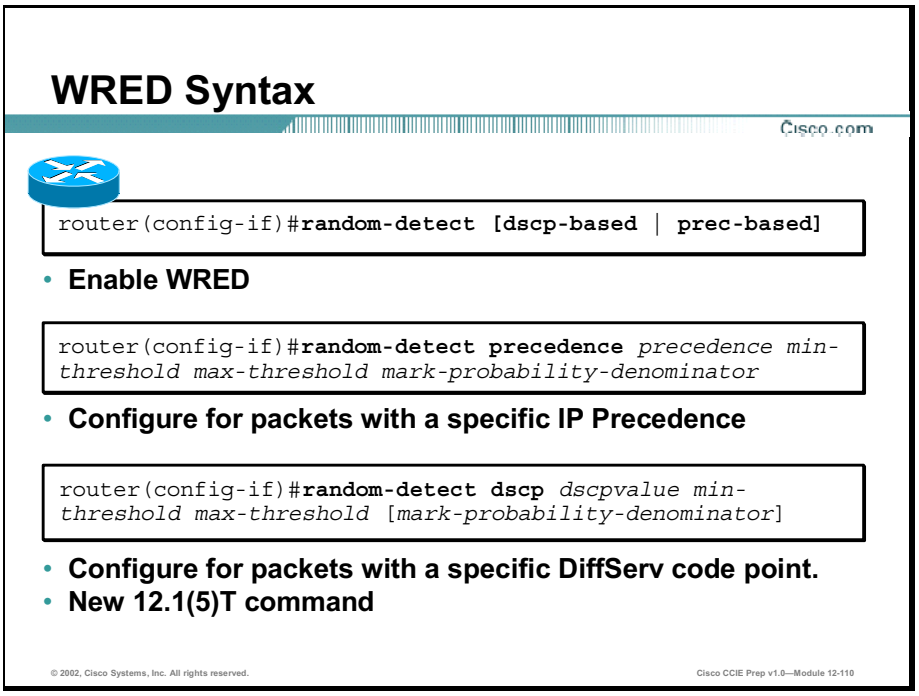

To enable Weighted Random Early Detection (WRED) on an interface, use the command **random-detect** [**dscp-based** | **prec-based**], where **dscp-based** will cause WRED to discard based on Differentiated Services Code Point values, and **prec-based** will cause WRED to discard based on IP Precedence values. If neither dscp-based nor prec-based is specified, then WRED will default to prec-based.

Even though WRED has default values for when it will begin discarding packets, and what that probability of discard is for various IP Precedence and DSCP values, these parameters may be altered. To specify the characteristics for **prec-based** WRED, use the command **random-detect precedence** ***precedence min-threshold max-threshold mark-probability-denominator***, where ***precedence*** is an IP Precedence value in the range of 0 – 7, ***min-threshold*** specifies the average queue depth after which WRED will begin discarding packets, ***max-threshold*** is the queue depth after which all packets will be discarded, and ***mark-probability-denominator*** is the fraction of packets dropped when the queue depth approaches the max-threshold. For example, a mark probability denominator of 100 would indicate that when the queue depth equaled the max-threshold, there would be a 1 in 100 chance that a router would discard a packet with the specified IP Precedence value. The probability of packet discard increases linearly from the min-threshold, with a zero probability, up to the max-threshold with a probability defined as 1/(*mark-probability-denominator*). Similarly, WRED's discard parameters can be adjusted for dscp-based WRED with the command **random-detect dscp** ***dscpvalue min-threshold max-threshold*** [***mark-probability-denominator***], where ***dscpvalue*** can be a DSCP value in the range of 0 – 63.

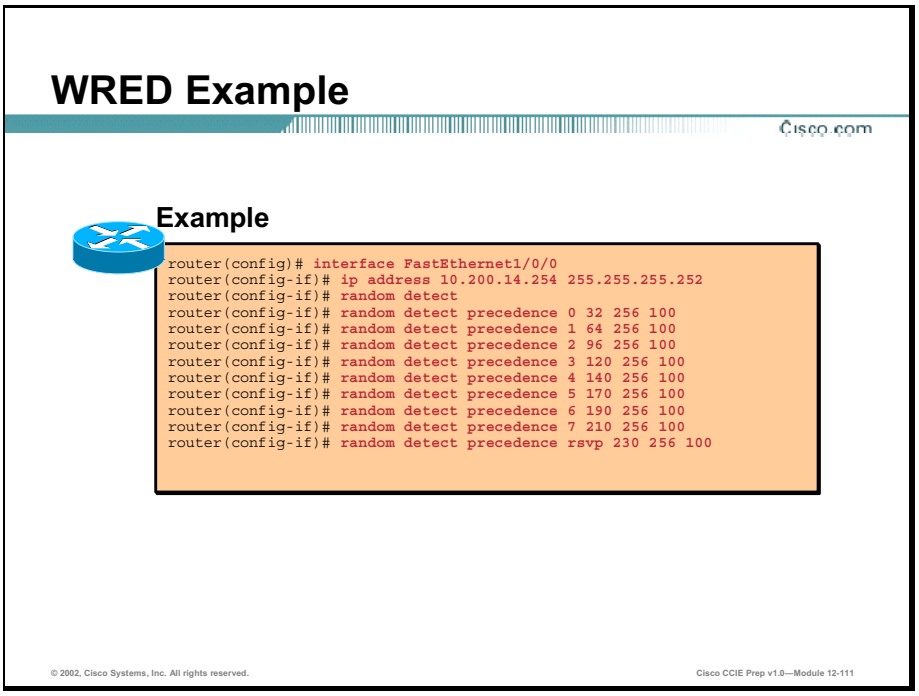

In the example shown, IP Precedence-based WRED is enabled on the FastEthernet 1/0/0 interface. Additionally, the minimum threshold, maximum threshold, and mark probability denominator values have been altered such that as any packet approaches the **max-threshold** queue depth, there will be a 1 in 100 chance that the packet will be discarded. Notice that higher priorities have higher **min-threshold** values, indicating that the queue depth would have to be greater for high priority packets to be discarded.

QoS Verification

This section examines the **show** commands that can be used to verify and troubleshoot Quality of Service configurations.

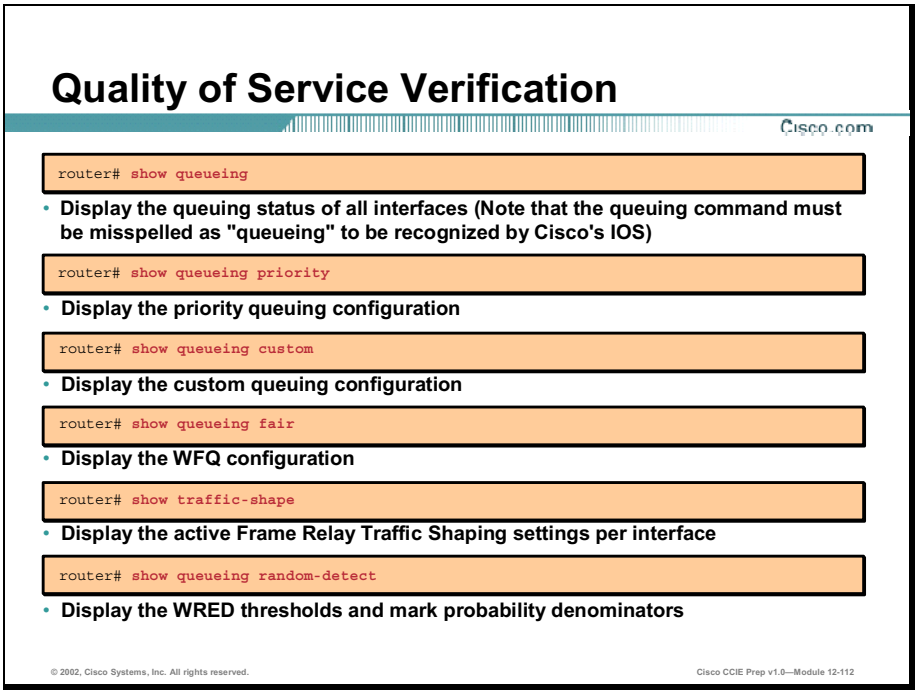

Cisco IOS provides several **show** commands for verification and troubleshooting of QoS settings. To view the queuing status of all interfaces, issue the command **show queueing**. Note that the queuing command must be misspelled as “queueing” to be recognized by the Cisco IOS.

To view the configuration for a specific queuing type, the following commands may be used: **show queueing priority**, **show queueing custom**, and **show queueing fair**, for PQ, CQ, and WFQ respectively.

Frame Relay Traffic Shaping (FRTS) settings may be examined by issuing the command **show traffic-shape**. The output will display such FRTS parameters as CIR, Bc, Be, and Tc.

WRED parameters may be viewed by entering the command **show queueing random-detect**. The output will show the minimum threshold, maximum threshold, and mark probability denominator settings for the various IP Precedence or DSCP values.

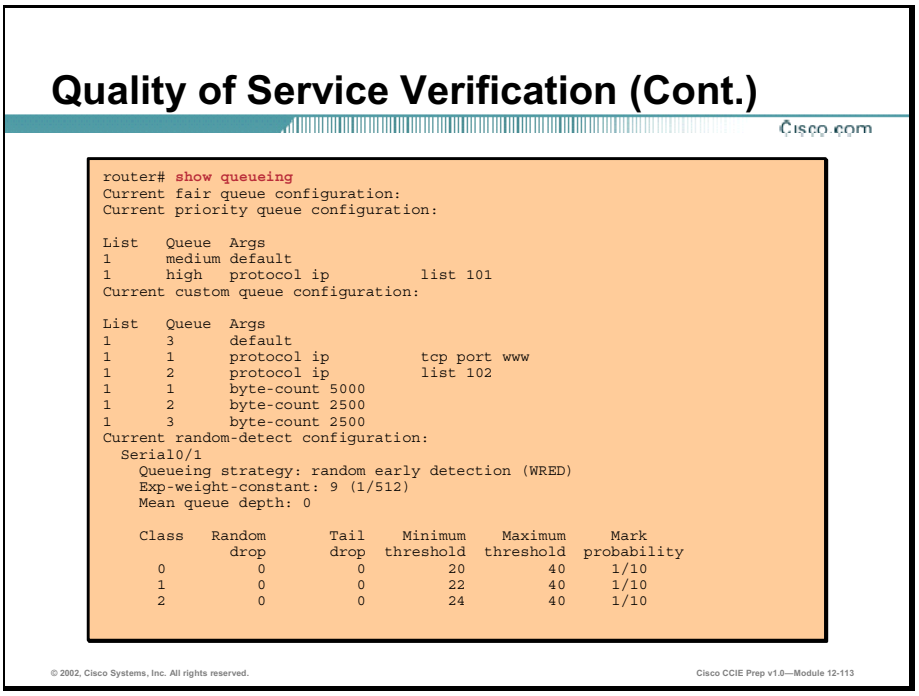

The output shown is from the **show queueing** command, which displays the queuing configuration for all interfaces. The output details the traffic types assigned to the predefined priority queues, the traffic types and the byte counts assigned to the custom queues, and the various WRED drop thresholds for each IP Precedence value.

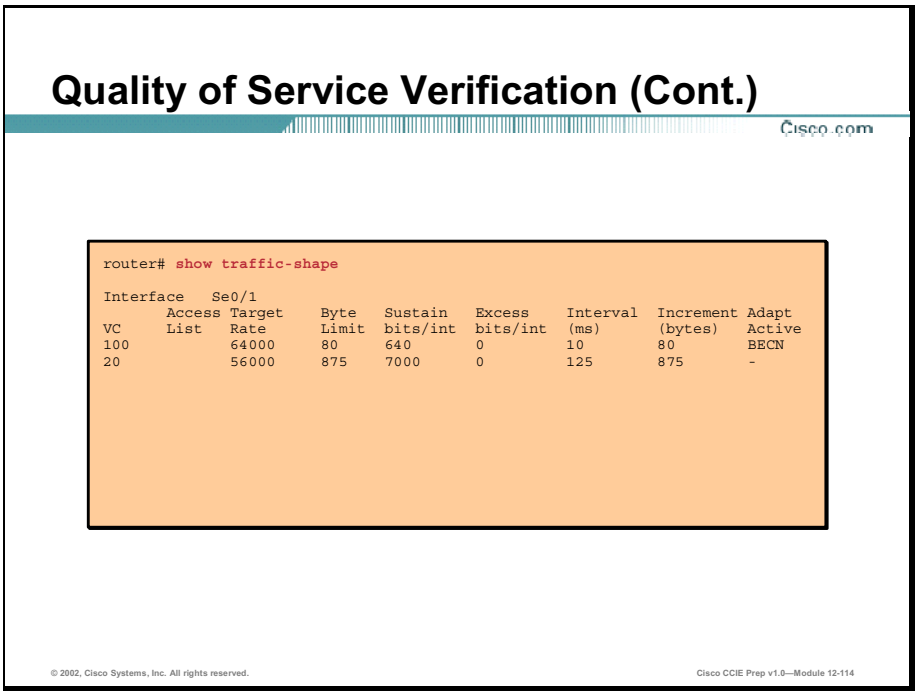

The output shown is from the **show traffic-shape** command, which displays the FRTS parameters, assigned to each DLCI. In this example, DLCI 100 has a CIR of 64 kbps, a Bc of 640 bits per time interval, a Be of 0 bits per time interval, a Tc of 10 ms, and the DLCI responds to BECNs.

DLCI 20 has a CIR of 56 kbps, a Bc of 7,000 bits per time interval, a Be of 0 bits per time interval, a Tc of 125 ms, and the DLCI does not respond to BECNs. Notice that both of the DLCIs are on the same physical interface, Serial 0/1.

Summary

This section summarizes the key points you learned in this lesson.

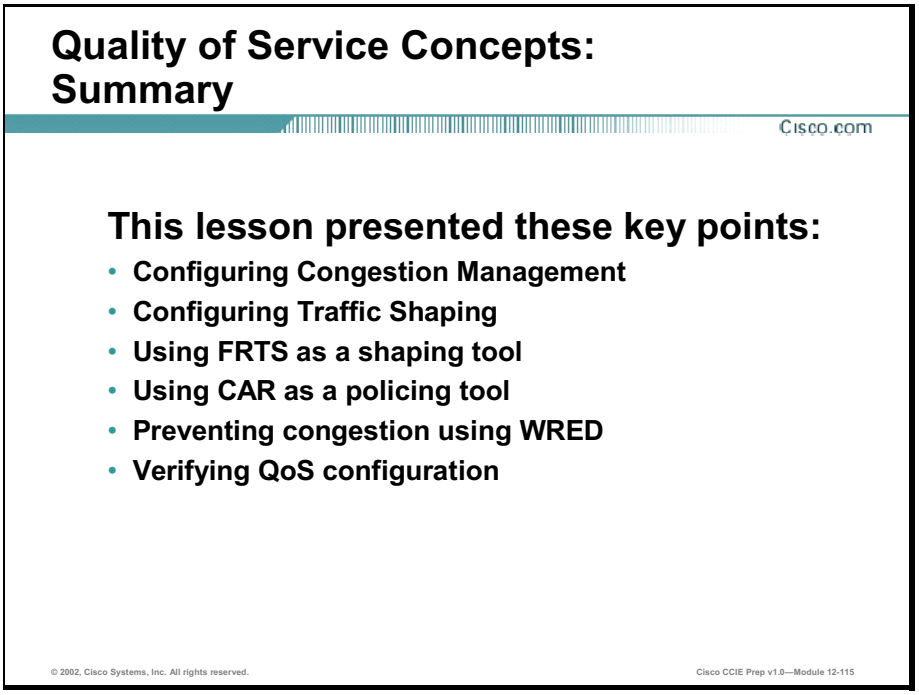

Next Steps

After completing this lesson, go to:

- Appendix

References

For additional information, refer to these resources:

- Deploying Quality of Service in the Enterprise (DQoS) Chapters 4, 5, and 7.
- QoS Configuration – http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121cgcr/qos_c/

Lesson Assessment (Quiz)

Q1) Which two of the following tools can limit the maximum amount of bandwidth sent over an interface?

A) CQ

B) PQ

C) WRED

D) CAR

E) WFQ

F) FRTS

Q2) If Frame Relay Traffic Shaping were configured with a Bc of 5600 bits per time interval and a CIR of 56 kbps, what would be the Tc?

A) 125 ms

B) 100 ms

C) 10 ms

D) 1.25 ms

Q3) Which command would configure WRED with a 10 percent discard probability for IP Precedence 3 packets at a queue depth of 256?

A) random-detect precedence 3 120 256 10

B) random-detect precedence 3 64 256 100

C) wred precedence 3 120 256 10

D) wred precedence 3 64 256 100

Q4) Which of the following is NOT a valid CAR conform or exceed action?

A) continue

B) drop

C) retransmit

D) set-prec-continue *new-prec*

E) transmit

F) set-prec-transmit *new-prec*

Q5) Which congestion management tool services queues in a round robin fashion?

A) CQ

B) PQ

C) WRED

D) CAR

E) WFQ

F) FRTS

A

Appendix A: Configuring a Terminal Server

Overview

This document describes the advantages of using terminal server and reverse telnet.

Terminal Server Advantages

A router with multiple asynchronous lines, such as a 2509 or 2511, can be used as a terminal server to provide remote access to other routers and switches via their console ports. There are many advantages of having remote console access to a device versus telnet access. One of those advantages is the ability to access a router remotely without any telnet configuration on the router. This enables you to remove the configuration on a router and still be able to access the router remotely. Another major advantage is the ability to remotely perform password recovery.

The 2509 and 2511 are the most common devices used for terminal servers. The 2509 provides eight asynchronous lines and the 2511 provides sixteen. If you are building a home lab, there are some less expensive alternatives, if you can obtain them, as they are no longer sold by Cisco. They are the cs-508 and cs-516. These devices can be used as lower end terminal servers.

Reverse Telnet

The terminal server provides remote console access to devices via a process known as reverse telnet. Reverse telnet allows you to telnet from a device to a certain line number on the device.

Here is a sample configuration of a 2509 configured as a terminal server.

```
hostname term_serv
!
no ip domain-lookup
ip host R1 2001 10.1.1.1
ip host R2 2002 10.1.1.1
ip host R3 2003 10.1.1.1
ip host R4 2004 10.1.1.1
ip host R5 2005 10.1.1.1
ip host R6 2006 10.1.1.1
ip host frame-switch 2007 10.1.1.1
ip host cat-switch 2008 10.1.1.1
!
interface Loopback0
 ip address 10.1.1.1 255.255.255.0
!
interface Ethernet0
 ip address 192.168.0.1 255.255.255.0
!
interface Serial0
 no ip address
 shutdown
!
interface Serial1
 no ip address
 shutdown
!
ip classless
!
line con 0
line 1 8
 no exec
 transport input all
line aux 0
line vty 0 4
 password cisco
 login
end
```

There are two key elements in this sample configuration that allow this router to act as a terminal server. The first key element is under the **line 1 8** configuration. Notice the two commands used here: **transport input all** and **no exec**. The first command **transport input all**

allows all protocols on lines 1-8 including telnet. This allows the router to perform reverse telnet.

The second command no exec prevents exec processes on these lines. This is recommended, as these lines will not be used to connect to the terminal server itself. This command prevents garbage text from the commands issued on other routers from appearing on the terminal server's console.

Line Numbering

In order to configure reverse telnet on a terminal server, you must understand the line numbering that Cisco routers follow. Any line on a Cisco router can be addressed by a corresponding port number. The port numbers for reverse telnet are 2000 + the line number. Therefore, port 2001 refers to line 1, 2002 refers to line 2, and so on.

After all asynchronous line numbers comes the line number for the aux port. On a 2509, the aux port would be line 9. That means that you can also reverse telnet to it via the port number of 2009, increasing the number of reverse telnet connections allowed on a 2509 to 9 instead of 8.

With line 1 plugged into the console port on router R1, R1 can be accessed by telnetting to any IP address on the terminal server with the port number 2001.

For example: telnet 10.1.1.1 2001

Often a loopback address is configured strictly for reverse telnet purposes, because loopbacks are virtual interfaces that never go down.

The second key element in this sample configuration is the ip host statements. The ip host statements are used to map a name to an IP address and port number. Therefore, instead of using the command **telnet 10.1.1.1 2001** to access R1, you could type **telnet R1**, or more

simply **R1**. The router assumes that you want to use telnet when an IP address or host name is typed at the command prompt.

Once you have successfully reverse telnetted to device, you need to be able to suspend telnet sessions and quickly move between them. To suspend a telnet session and get back to the terminal server, use the key combination **Ctrl+Shift+6** then **X**. You can now telnet to another device and suspend that session. You will be able to quickly switch between suspended telnet sessions by using the session number. The current sessions can be viewed with the **show sessions** command.

```
termsrv#show session
Conn Host                Address             Byte  Idle Conn Name
*  1 R1                  100.1.1.1              0           0      R1


termsrv#1
[Resuming connection 1 to R1 ... ]
```

For example, to switch to the telnet session to R1, you simply enter **1** and press **Enter**.

B

Appendix B: Configuring a Frame Relay Switch

Overview

This section shows you how to set up a router as a Frame Relay switch.

Configuring the Router as a Frame Relay Switch

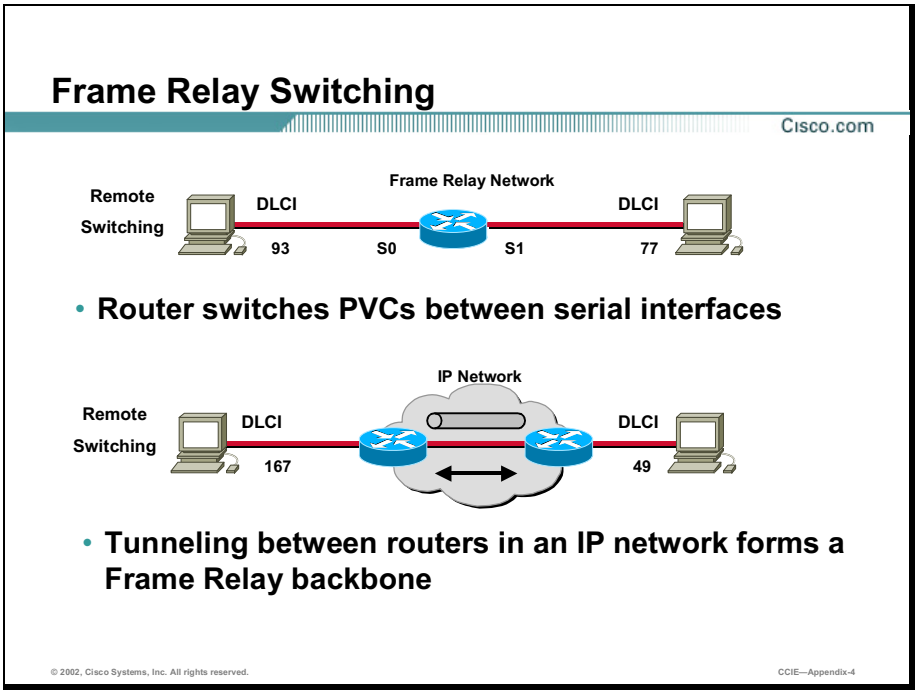

Local Frame Relay switching enables the Cisco router to switch Frame Relay frames between interfaces based on the data-link connection identifier (DLCI) number in the frame header. A router interface performing PVC switching is usually configured as a Frame Relay switch.

Remote Frame Relay switching enables the router to encapsulate Frame Relay frames in IP datagrams and tunnel them across an IP backbone. The Cisco generic routing encapsulation (GRE) tunnel protocol is used for remote Frame Relay switching. The router is usually configured as a Frame Relay DCE device.

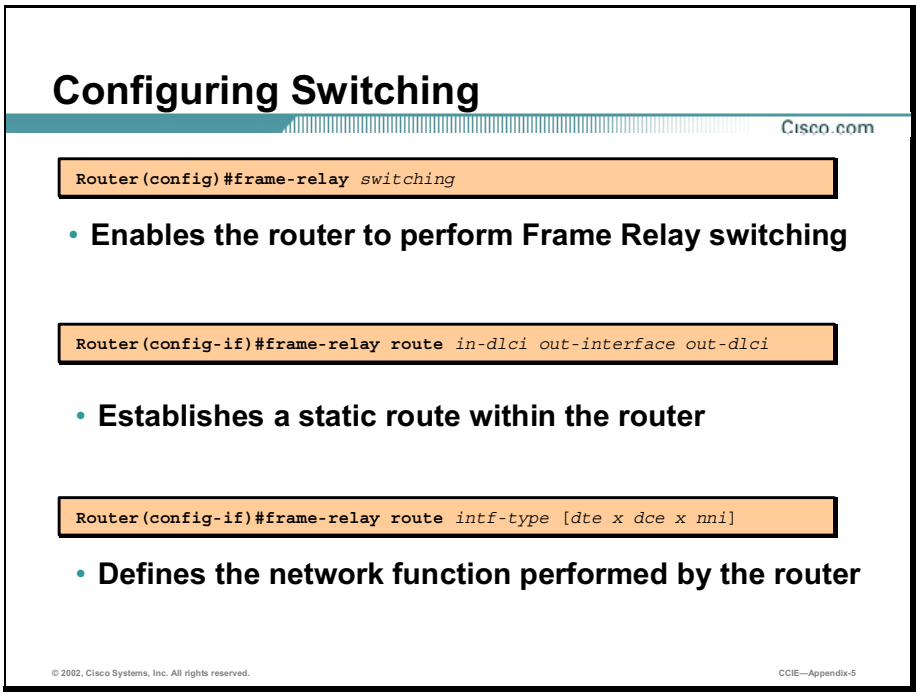

Use the **frame-relay route** command to link traffic inside the router between two serial ports when the router is functioning as a Frame Relay switch. The router performs PVC switching between the serial ports.

Table 1: frame-relay route *in-dlci out-interface out-dlci* Command

Command	Description
`in-dlci`	DLCI on which the packet is received on the interface.
`out-interface`	Interface the router uses to transmit the packet.
`out-dlci`	DLCI the router uses to transmit the packet over the specified out-interface.

Table 2: frame-relay intf-type Command

Command	Description
`dte`	(Optional) Router is connected to a Frame Relay network.
`dce`	(Optional) Router is connected to another router and is acting as a Frame Relay switch.
`nni`	(Optional) Router functions as a Frame Relay switch and is connected to another switch performing Network-to-Network Interface (NNI) support.

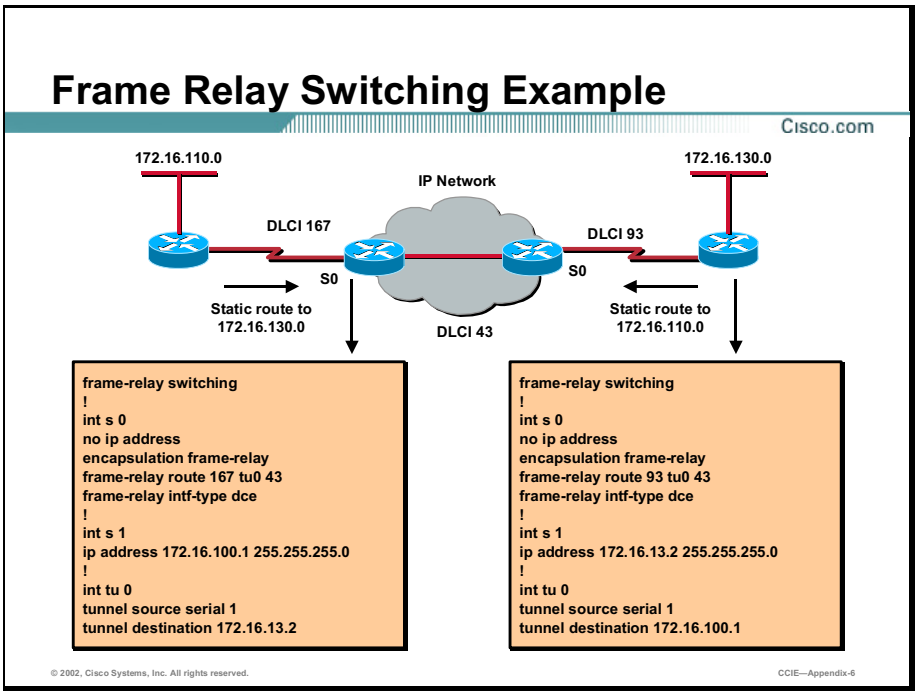

Use the **frame-relay intf-type** command to configure the interface to function as a Frame Relay switch. The type of Frame Relay switch is determined by the router's function within the Frame Relay network.

The example uses the following commands:

Table 3: frame-relay route 167 tu0 43 Command

Command	Description
`167`	Specifies the DLCI of the arriving (source) traffic to be switched.
`tu0`	Specifies the outgoing interface to use.
`43`	Specifies the outgoing DLCI to use when forwarding the traffic.
`frame-relay intf-type dce`	Establishes interface S0 as the DCE. In this back-to-back Frame Relay connection, one interface must act as the DCE.
`tunnel source serial 1`	Defines that software-only tunnel interface 0 will use physical interface serial 1 as the entry into the tunnel.
`tunnel destination 172.16.13.2`	Defines that the tunnel will deliver traffic to IP address 172.16.13.2 as the tunnel destination.

The router is configured as a remote Frame Relay switch. Traffic arriving on S0 using DLCI 167 will be switched to output interface S1 and DLCI 43 will be used in the source identifier. The traffic will be carried through the IP network using a GRE tunnel having a next-hop destination of 172.16.100.1.

The tunnel uses the same DLCI number.

Complete Frame Relay Switch Configuration

Following is the configuration output from a router that is acting as a Frame Relay switch. In this example, a 3600 was configured as a Frame Relay switch.

```
3640-switch#sh run
Building configuration...
Current configuration:
!
version 11.3
no service password-encryption
!
hostname 3640-switch
!
enable password cisco
!
frame-relay switching
!
! <output omitted>
!
interface Serial1/0
no ip address
encapsulation frame-relay
clockrate 64000
frame-relay intf-type dce
frame-relay route 110 interface Serial1/1 100
!
interface Serial1/1
encapsulation frame-relay
clockrate 64000
frame-relay intf-type dce
frame-relay route 100 interface Serial1/0 110
! <Output Omitted>
!
!
ip classless
!
!
line con 0
line aux 0
line vty 0 4
login
!
end
```

Enables router as switch.

Enables Frame Relay on interface.

DCE sets

Instructs router to act as DCE and forward LMI information.

Route between each router. This is the information that is carried in the LMI Status enquiry frame sent to the DTE routers.

C

Appendix C: Configuration Register Settings

Overview

This document describes the 16-bit boot register.

Gaining Privileged Access: The 16-Bit Boot Register

One of the best-kept secrets of Cisco routers and switches is the 16-bit boot register. The 16-bit register is located on almost every Cisco platform in one variation or another. For example, this is the same register that was set by jumpers on the AGS series routers in the early 1990s. It is the same register that is found in the Catalyst switches in 2001. And, for the most part, it is the same register on all Cisco routers, sometimes masked in a utility called CONFREG.

Another common example of using the boot register is during password recovery. The boot register, actually bit 6, is the bit that you flip when you change the register from 0x2102 to Ox2142 during password recovery. During password recovery, bit 6 is set to ignore NVRAM on startup. This is perhaps the most common use of the register. Some other uses of the boot register include the following:

- Recovering a lost password
- Enabling or disabling the console Break key
- Allowing manual boot of the OS using the B command at the bootstrap program (ROM monitor) prompt
- Changing the router boot configuration to allow a Flash or ROM boot
- Performing maintenance testing from the ROM monitor

- Loading an image into Flash memory
- Permanently disabling a router

Because the boot register represents the "keys" to your router, it is important to explain the entire register rather then covering just bit 6.

To display the boot register, key in the show version command. The boot register is displayed at the bottom of the text. Example C-1 demonstrates the show version command.

Example C-1 The show version Command, with a Boot Register Set to Boot to ROM, Ox2l0l

```
router(boot)#show version

Cisco Internetwork Operating System Software

IOS (tm) 3000 Bootstrap Software (IGS-RXBOOr), Version 10.2(8a), RELEASE SOFTWAR
E (fc1)

Copyright (c) 1986-1995 by cisco Systems, Inc. Compiled rue 24-Oct-95 15:46 by
mkamson

Image text-base: 0x01020000, data-base: 0x00001000

ROM: System Bootstrap, Version 5.2(8a), RELEASE SOFTWARE

router uptime is 34 minutes System restarted by power-on Running default
software

cisco 2500 (68030) processor (revision L) with 14332K/2048K bytes of memory.
Processor board serial number 03071163 with hardware revision 00000000 X.25
software, Version 2.0, NET2, 8FE and GOSIP compliant. ISDN software, Version
1.0.

1 Ethernet/IEEE 802.3 interface. 2 Serial network interfaces. 1 ISDN 8asic Rate
interface.

32K bytes of non-volatile configuration memory.

16384K bytes of processor board System flash (Read/Write)

Configuration register is 0x2101 router(boot)#
```

The boot register is formatted with the most-significant bit on the right, as illustrated by Figure C-1. This figure also shows how the default settings of 0x2102 are derived on Cisco routers.

Figure C-1 Default Settings of the 16-Bit Boot Register

Bit 15	Bit 14	Bit 13	Bit 12	Bit 11	Bit 10	Bit 9	Bit 8	Bit 7	Bit 6	Bit 5	Bit 4	Bit 3	Bit 2	Bit 1	Bit 0
0	0	1	0	0	0	0	1	0	0	0	0	0	0	1	0
2				1				0				2			

Briefly stepping through the default settings of the register, you can see that bits 1,8, and 13 are set to 1, or the ON position. Having bit 1 set then sets the boot portion of the register to a hexadecimal value of 2. This tells the router to boot from Flash if a valid IOS is found there. Having bits 4 through 7 set to 0 enables the router to boot normally; from NVRAM, preserve the banner and set "all 1s" as the broadcast. Bit 8 tells the router that the Break key is disabled. The rest of the register sets the network broadcast to 1 s, sets the console baud rate to 9600, and

determines how the router responds to a netboot failure. As mentioned previously, the most common use of this register is the flipping of bit 6, causing the router to ignore the startup config stored in NVRAM. Again, this is the same procedure used in password recovery.

Table C-1 illustrates the entire register and its settings in detail. Refer to this table when reading the following detailed descriptions of the boot register.

Table C-1 The Entire 16-Bit Boot Register with Default Settings

Bit	Meaning	Default Setting
0-3	Boot Field: 0x0= Boot ROM monitor. --- 0x1 = Boot from onboard ROM, or boot to boot mode, if a subset of the IOS exists. --- 0x2 to 0xF Causes the following (listed in order of precedence): Boot from Flash, if a valid IOS file exists. Follow **boot system** commands found in the configuration. Use the register value to form a filename from which to netboot a system image from.	0 0 I 0
4	Fast boot: Force load through the **boot system** commands found in the configuration.	0
5	High-speed console: 1 = console operates at 19.2 or 38.4; works with bits 11 and 12.	0
6	Ignore startup-config file: 1 = ignore NVRAM.	0
7	OEM bit: 1 = disabling the display of the Cisco banner on startup.	0
8	Break key: 1 = disable.	1
9	Not used.	0
10	Netboot broadcast format: Setting bit 10 = 1 causes the processor to use an all-zeros broadcast.	0

continued

Bit	Meaning	Default Setting
11-12	Console baud rate: **Bit 5** = 1 **Bit 11** = 1 **Bit 12** = 0 **Console baud rate** = 38,400 --- **Bit 5** = 1 **Bit 11** = 0 **Bit 12** = 0 **Console baud rate** = 19,200 --- **Bit 5** = 0 **Bit 11** = 0 **Bit 12** = 0 **Console baud rate** = 9600 --- **Bit 5** = 0 **Bit 11** = 0 **Bit 12** = 1 **Console baud rate** = 4800 --- **Bit 5** = 0 **Bit 11** = 1 **Bit 12** = 1 **Console baud rate** = 2400 --- **Bit 5** = 0 **Bit 11** = 1 **Bit 12** = 0 **Console baud rate** = 1200	00 00
13	Response to netboot failure: 1 = boot from ROM after netboot 1 failure, 0 = continue to netboot.	1
14	Netboot subnet broadcast: Setting bit 14 = 1 forces a subnet broadcast.	0
15	Enable diagnostic messages: 1 = ignore NVRAM and display diagnostic messages.	0

Boot Field (Bits O Through 3)

The boot field controls the booting of the router. This field starts with the first 4 bits on the right. If this field is set for 0x0, decimal 0, the router will boot to ROM monitor mode. For example, setting the register for 0x2100 causes the router to boot to ROM monitor mode. Setting this value to 0x1 causes the router to boot from its onboard ROM. This ROM may contain a full IOS, such as in the 7000 series, or a subset of the IOS, as in the 2500 series. The prompt, when in boot mode, is represented with (boot) behind the router's host name.

If you set the boot field to a value of 2 through F, and if there is a valid system boot command stored in the configuration file, the router boots the system software as directed by that value. If you set the boot field to any other bit pattern, the router uses the resulting number to form a default boot filename for netbooting. The router creates a default boot filename as part of the automatic configuration processes. To form the boot filename, the router starts with cisco and links the octal equivalent of the boot filename, a dash, and the processor-type name. A Cisco 4000 with the bit pattern of 0x1 set in the first octet will try to load a TFTP file named Cisco2-4000. Table C-2 lists the default boot filenames or actions for the processor when setting the boot field bits. The xxxx stands for the processor type – for instance, in Cisco 4000, xxxx = 4000.

Table C-2 Default Boot Filenames

Action/Filename	Bit 3	Bit 2	Bit 1	Bit 0
Boot to ROM monitor	0	0	0	0
Boot from ROM	0	0	0	1
cisco2-xxxx	0	0	1	0
cisco3-xxxx	0	0	1	1
cisc04-xxxx	0	1	0	0
cisco5-xxxx	0	1	0	1
cisc06-xxxx	0	1	1	0
cisco7-xxxx	0	1	1	1
cisco10-xxxx	1	0	0	0
cisco11-xxxx	1	0	0	1
cisco12-xxxx	1	0	1	0
cisco13-xxxx	1	0	1	1
cisco14-xxxx	1	1	0	0
cisco15-xxxx	1	1	0	1
cisco16-xxxx	1	1	1	0
cisco17-xxxx	1	1	1	1

Fast Boot/Force Boot (Bit 4)

Setting this bit forces the router to load the Cisco IOS Software found in the configuration set by the **boot system flash** command. If no Cisco IOS Software matches the filename set by this command, the router will boot to boot mode. For example, adding the line **boot system flash c2500-js56-1.120-3.bin** forces the router to look for the file c2500-js56-1.120-3.bin in Flash memory. If an exact match of this filename is not found, the router will boot in boot mode.

High-Speed Console (Bit 5)

The setting of bit 5 works in conjunction with bits 11 and 12. Setting this bit is for high-speed console access above 9600 bps. When this bit is set, you can connect to the console port at speeds of 19,200 bps and 38,400 bps. For a complete listing of how the jumper works in conjunction with bits 10, and 11, see Table C-4.

Caution Bit 5 is an "undocumented" bit for a reason. The console port is critical to router operation and troubleshooting. The higher the data speeds are, the more sensitive the connection is and the higher the probability is that you will not be capable of connecting to the router at these high speeds. If you do not have Telnet access or another "back door" into the router enabled, the consequences can be dire. The gains from operating the console port at 19,200 bps or 38,400 bps instead of 9600 bps are minor. Keep in mind that the uses for this interface are for router key-ins and configuration; it is not necessary to have high-speed console access. Change this bit with extreme caution.

Ignore NVRAM (Bit 6)

Setting this bit forces the router to ignore the configuration file in NVRAM, called the *startup-config.* When you ignore NVRAM, you essentially are ignoring the startup-config. You can still view the startup-config with the **show** command, but the configuration will be absent from the running-config. This is also the bit that is flipped during password recovery.

GEM Bit (Bit 7)

This bit was created for Original Equipment Manufacturers (OEMs) versions of the routers. By setting this bit, the Cisco Systems, Inc. banner will be ignored. If the IOS has encryption software on it, the encryption warning will still be displayed.

Break Key (Bit 8)

Setting this bit disables the Break key. If you set this bit to 0, then at any time during the routers uptime – not just during the boot process – you can halt the operating system with the press of a single key. This is a powerful setting and should not be changed. Disabling the break – it is disabled by default – does not affect the Break key during the first 60 seconds of initialization. During this time, the Break key will still halt the router.

Reserved (Bit 9)

This bit is currently not in use.

Netboot Broadcast Format (Bits 10 and 14)

Setting bits 10 and 14 controls how the routers and switches handle subnet and host broadcasts. The default broadcast address is all 1s in the host or subnet destination address. Changing these bits allows for backward compatibility for many older UNIX hosts, such as Berkley UNIX 4.2BSD. Most IP implementation today uses a 1s compliment for broadcast messages, so you probably will never modify these settings. Table C-3 illustrates the use of bit 10 and bit 14.

Table C-3 Configuration Settings for Broadcast Address Control, Bit 10 and Bit 14

Bit 14	Bit 10	Address (<net><host>)
0	0	<1s> <1s>
0	1	<0s> <1s>
1	0	<net> <1s>
1	1	<net> <0s>

System Console Terminal Baud Rate Settings (Bits 5, 11, and 12)

Bits 5, 11, and 12 control the baud rate (bps) of the console port. The routers are shipped with this setting to 9600, which has bits 5, 11, and 12 off, or set at 0. Table C-4 shows the baud rate settings. For example, to increase the baud settings of the routers console port, use a register of 0x2122 for 19.2 access.

Table C-4 Configuration Settings for System Console Baud Rate

Bit 5	Bit 11	Bit 12	Console Baud Rate
1	1	0	38,400 bps
1	0	0	19,200 bps
0	0	0	9600 bps
0	0	1	4800 bps
0	1	0	1200 bps
0	1	1	2400 bps

Netboot Failure Response (Bit 13)

Setting bit 13 causes the router to load the Cisco IOS Software from the default location after five netboot failures. The default for this bit is on, or 1, which is why most of the routers' jump registers start with 2. Setting this bit to 0 causes the router to continue to netboot and never look at the ROM for booting.

Display Factory Diagnostics (Bit 15)

Setting bit 15 causes the router to display factory diagnostic messages. Setting this bit also forces NVRAM to be ignored. To display these diagnostic messages, configure the register at 0xA102. The A sets bit 15 and bit 13, forcing diagnostics messages to appear during initialization.

Understanding the Boot Process

This next section can be found in a similar format on the Cisco documentation CD that comes with all new Cisco routers. Although everything can be found on the CD, this section is important enough to highlight:

When a router is powered on or rebooted, the following events happen:

- The ROM monitor initializes
- The ROM monitor checks the configuration register boot field (the lowest 4 bits in the register.)
  - If the boot field is 0x0, the system does not boot an IOS image and waits for user intervention at the ROM monitor prompt
  - If the boot field is 0xl, the ROM monitor boots the boot helper image. (On some platforms the boot helper image is specified by the BOOTLDR environment variable.)
  - If the boot field is 0x2 through 0xF, the ROM monitor boots the first valid image specified in the configuration file or specified by the BOOT environment variable
- When the boot field is 0x2 through 0xF, the router goes through each command in order until it boots a valid image. If bit 13 in the configuration register is set, each command will be tried once. If bit 13 is not set, the **Boot system** command specifying a network server will be tried up to five more times. The timeouts between each consecutive attempt are 2, 4, 16, 256, and finally 300 seconds. If it cannot find a valid image, the following events happen:
  - If all boot commands in the system configuration file specify booting from a network server and all commands fail, the system attempts to boot the first valid file in Flash memory.
  - If the boot-default-ROM-software option in the configuration register is set, the router will start the boot image (the image contained in boot ROM or specified by the BOOTLDR environment variable).
  - If the boot-default-ROM-software option in the configuration register is not set, the system waits for user intervention at the ROM monitor prompt. You must boot the router manually.
  - If a fully functional system image is not found, the router will not function and must be reconfigured through a direct console port connection.

- When looking for a bootable file in Flash memory:
  - The system searches for the filename in Flash memory. If a filename is not specified, the software searches through the entire Flash directory for a bootable file instead of picking only the first file.
  - The system attempts to recognize the file in Flash memory. If the file is recognized, the software decides whether it is bootable by performing the following checks:
    - For run-from-Flash images, the software determines whether it is loaded at the correct execution address.
    - For run-from-RAM images, the software determines whether the system has enough RAM to execute the image.

This process changes on platforms with dual processor cards or dual Flash cards, such as those that are found on the 7000 series or in the Catalyst RSM. Figure C-2 diagrams this rather complicated process as it is found on most platforms (except those noted).

Accessing the Register

The boot register is a 16-bit register represented in hex to the router. The router make and model determine how the register is accessed. As mentioned previously, the AGS used 16 jumpers to set this register. Every router and switch allows access to the register through the configuration, assuming that you have privileged-level access. Switches work much in the same way as routers. First, you will learn about accessing the register on Catalyst switches, and then you will learn about routers.

Figure C-2 Router Boot Process

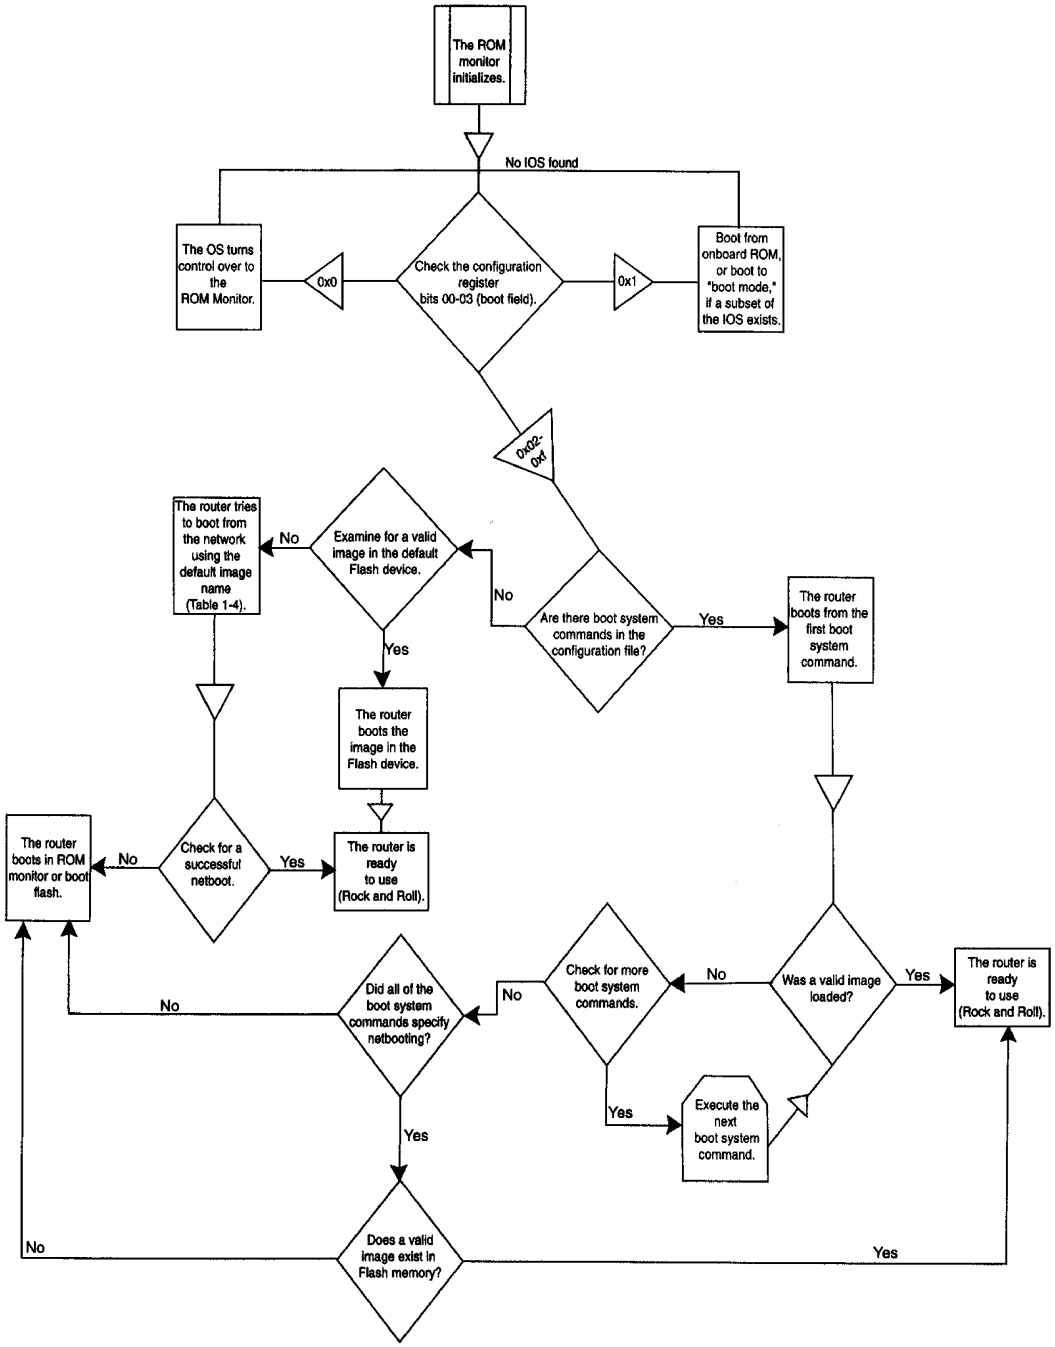

Accessing and Configuring the Register: Catalyst Switches

For the most part, the 16-bit register is identical to its cousin found in the router. The differences are slight. Most of the bits that are used in netbooting are used for broadcast control and are not used on the Catalyst switches. Bit 6 operates differently on the Catalyst than it does on the router. Setting Bit 6 clears the configs from NVRAM, which is the same as entering the **clear config aII** command – that is, it clears the entire configuration stored in NVRAM the next time the switch is restarted.

The initialization process on the Catalyst 5000 series Supervisor Engine III and the Catalyst 4000, 2948G, and 2926 series switches involves two software images: the ROM monitor and the supervisor engine system code. When the switch is reset, the ROM monitor code is executed first. Then, depending on the boot register settings in NVRAM, the switch either remains in ROM monitor mode or loads the supervisor system image. If a fatal exception error occurs during power up, the switch remains in ROM monitor mode. Figure 1-11 illustrates the 16-bit boot register for the Catalyst series of switches. Table 1- 7 provides detailed descriptions of the boot register.

Figure C-3 The Entire 16-Bit Boot Register for Catalyst Switches with Default Settings

Bit 15	Bit 14	Bit 13	Bit 12	Bit 11	Bit 10	Bit 9	Bit 8	Bit 7	Bit 6	Bit 5	Bit 4	Bit 3	Bit 2	Bit 1	Bit 0
0	0	1	0	0	0	0	1	0	0	0	0	0	0	1	0
0				1				0				F			

Table C-4 Catalyst Switch Boot Register Bit Meanings and Default Settings

Bit	Meaning	Default Setting
0-3	Boot Field: 0x0= Boot ROM monitor. 0x1 = Boot from onboard ROM, or boot to boot mode, if a subset of the IOS exists. 0x2 to 0xF Causes the following (listed in order of precedence): Boot from Flash, if a valid IOS file exists. Follow **boot system** commands found in the configuration. If a boot image in the BOOT environment variable list is not found, boot in ROM monitor mode.	1111
4	Reserved	0
5	Reserved	0
6	Clear NVRAM: 1 = Clear NVRAM.	0
7	OEM bit: 1 = disabling the display of the Cisco banner on startup (Not used.)	0
8	Break key: 1 = disable.	1
9	Unsupported baud rate.	0
10	IP will use an all-zeros broadcast. (Not used.)	0
11-12	Console baud rate: 00 = 9600 01 = 4800 10 = 1200 11 = 1200 On the Catalyst 4000 and 2948G, this speed is fixed at 9600	00
13	Boots default Flash if network boot fails. (Not used.)	0
14	Netboot subnet broadcast: Setting bit 14 = 1 forces a subnet broadcast. (Not used.)	0
15	Enable diagnostic messages: 1 = ignore NVRAM and display diagnostic messages. (Not used.)	0

The default register is set for 0x010f. This allows the system to boot from the image specified in the BOOT environment variable; the console will operate at 9600 baud, and any configuration in NVRAM will be loaded. To display the current register settings, use the **show boot** [*module_number*] command. Example C-2 shows how to display the current configuration register and BOOT environment settings.

Example C-2 Demonstration of the show boot Command

```
Console>(enable) show boot
BOOT variable = slot0:cat5000-sup3.4-2-1.bin,1;bootflash:cat5000-sup3.3-2-
1b.bin,1;bootflash:cat5000-sup3.4-1-2.bin,1;

Configuration register is 0x10f
Ignore-config: disabled
Console baud: 9600
Boot: image specified by the boot system commands

Console>(enable)
```

The following is list of register-specific commands for the Catalyst family of switches:

- **set boot config-register 0x** *value* [*mode_num*]

This command directly configures the boot register at the bit level. This command affects all the bits in the register by modifying the entire boot register at once.

- **set boot config-register baud {1200 | 12400 | 14800 | 19600}**[*module_number*]

This configures the ROM monitor console port baud rate. The ROM monitor uses the baud rate specified in the configuration register only if it is different from the baud rate specified by the **set system baud** command.

- **set boot config-register ignore-config enable**

This command clears the entire configuration stored in NVRAM the next time the switch is restarted. This is essentially the same as using the **clear config all** command, followed by a reload.

- **set boot config-register boot** {**rommon | bootflash | system**} [*module_number*]

This command determines what boot method the switch will use during the next startup:

  - **rommon** = Boot to the ROM monitor
  - **bootflash** = Boot from the first image stored in the onboard Flash
  - **system** = Boot from the image specified in the BOOT environment variable.This is the default setting.

- **set boot system flash** *device*:[*filename*] [**prepend**] [*module_number*]

This command specifies an image to add to the BOOT environment variable. This also specifies what device that image exists on.

- **clear boot system flash** *device*:[*filename*][*module_number*]

This command clears a specific image from the BOOT environment variable.

- **clear boot system all**[*module_number*]

This command clears the entire BOOT environment variable.

Accessing and Configuring the Register: Cisco Routers

To set the register by the configuration mode, enter **config-register** ***<0x0000-0xFFFF>***. Example C-3 demonstrates how to change the configuration register from 2102 to 2142. This forces the router to ignore NVRAM during its initialization. To see if the configuration settings have taken effect, perform the **show version** command after changing the register.

Note	You should always check and document the current configuration register setting before changing it. This might come in handy if you have problems.

Example C-3 Changing the Boot Register Through the Configuration

```
Documenting the current setting

router#

router#show version

Cisco Internetwork Operating System Software

IOS (tm) 2500 Software (C2500-JS56-L), Version 12.0(3), RELEASE SOFTWARE (fc1)
Copyright (c) 1986-1999 by cisco Systems, Inc.

*** text omitted ***

32K bytes of non-volatile configuration memory.

16384K bytes of processor board System flash (Read ONLY)

Configuration register is 0x2102

router#
```

```
Change the setting to 0x2142.

router#conf t

Enter configuration commands, one per line. End with CNTL/Z.
router(config)#config-register 0x2142

router(config)#^Z

router#

router#show version

Cisco Internetwork Operating System Software

IOS (tm) 2500 Software (C2500-JS56-L), Version 12.0(3), RELEASE SOFTWARE (fc1)
Copyright (c) 1986-1999 by cisco Systems, Inc.

*** text omitted ***

32K bytes of non-volatile configuration memory.

16384K bytes of processor board System flash (Read ONLY)

Configuration register is 0x2102 (will be 0x2142 at next reload)
```


Note Whenever you change the boot register from the configuration mode, you are prompted to save your configuration before you reload the router. This prompt is generated from entering the configuration mode and exiting, regardless of any changes made to the configuration. The register setting is not part of the startup-config or running-config, so it is not necessary to save the configuration for the new jump register setting to take place.

Accessing and Configuring the Register: ROM Monitor

If you cannot access the router's configuration, such as in a password-recovery situation, you can force the Cisco IOS Software to halt and go into ROM monitor mode. To enter ROM monitor mode, you must send a break signal to the router. By default, the Break key is disabled by the boot register; consequently, a restart of the router is needed. Almost all Cisco routers and switches can be interrupted by sending the break signal during the first 60 seconds of initialization. There are many ways to send the break signal and to interrupt router and switch operations, the most common of which are documented in Table C-4.

Table C-4 Standard Break Key Combinations

Terminal-Emulation Software	Platform	Operating System	Key Combination
Hyperterm (Version 595160)	IBM-compatible	Windows 9x	Ctrl-F6-Break
Kermit	Sun workstation	Solaris	Ctrl-\L
Kermit	Sun workstation	Solaris	Ctrl-\B
MicroPhone Pro	IBM-compatible	Windows 9.x	Ctrl-Break
Minicom	IBM-compatible	Linux	Ctrl-A-F
ProCommPlus	IBM-compatible	DOS or Windows	Alt-B
Telix	IBM-compatible	DOS	Ctrl-End
Telnet to Cisco	IBM-compatible	--	Ctrl-]
Teraterm	IBM-compatible	Windows 9.x	Alt-B
Hyperterm	IBM-compatible	Windows 9.x	Break
Hyperterm	IBM-compatible	Windows 9.x	Ctrl-Break
Tip	Sun workstation	Solaris	Ctrl-], then Break or Ctrl-C
			~#
VT 100 Emulation	Data general	N/A	Fl6
Hypterm	IBM-compatible	Windows NT	Shift-6 Shift-4 Shift-B (^$B)
Z- TERMINAL	Mac	Apple	Command-B
--	Break-Out Box	--	Connect pin 2 (X-mit) to +V for half a second
--	Cisco to aux port	--	Control-Shift-6, then B
--	IBM-compatible	--	Ctrl-Break

If your portable or laptop computer is using Windows 95/98/2000 with HyperTerm, the break signal is usually issued by pressing the Function key and the Break key, sometimes located on the Page Down or Pause key.

On a full-size 101 keyboard with Windows 95/98 with HyperTerm, the break signal is issued by pressing the Ctrl-Break/Pause key.

On Windows NT, you must configure NT to send the break signal with a function key. Set the break by entering the characters **^$B** (**Shift 6**, **Shift 4**, and uppercase **B**). HyperTerm 5.0 private edition sends the break for the Windows NT platform without any additional configuration.

To access the register of a Catalyst 5000 or 2926G series switch, you can enter ROM monitor mode by restarting the switch and then pressing the **Break** key during the first 60 seconds of initialization. On the Catalyst 4000 and 2948G series switches, you can enter ROM monitor mode by restarting the switch and then pressing **Control-C** during the first five seconds of initialization.

When using any other terminal-emulation software, consult the manufacturer's instructions on sending a break signal.

When you have successfully sent the break signal, the router prompt will change to a > character or a **rommon *x*** > prompt. There are two prompts because there are two types of ROM monitors. One is built around the earlier 2000 series boards. It requires more of a manual manipulation of the boot registers. The other type of ROM monitor is built around the newer 3600 and RISC-based platforms. This ROM monitor uses a utility called CONFREG to manipulate the boot register. Table C-5 lists some common router types and the type of ROM monitor used. The easiest way to tell what type of ROM monitor is used in your router is to simply key in the ? for help. If the CONFREG utility appears, execute it by typing in **CONFREG**.

Table C-5 ROM Monitor Compatibility Matrix

CONFREG ROM Monitor	Basic ROM Monitor
Cisco 1003 series	Cisco 2000 series
Cisco 1600 series	Cisco 2500 series
Cisco 3600 series	Cisco 3000 series
Cisco 4500 series	Cisco 4000 series with 680x0
Cisco 7200 series	Cisco 7000 series 10.0 ROM
Cisco 7500 series	Cisco IGS series running IOS 9.1 in ROM
IDT Orion-based router	
AS5200 and AS5300 platforms	

First, you will learn about the Basic ROM monitor, and then you will learn about the utility called CONFREG. When you have successfully transmitted a break signal, you should get a screen that resembles Example C-4; also note the **Abort at** message.

Example C-4 Example of a Successful Break into ROM Monitor, Followed by the h or Help Command

```
System Bootstrap, Version 5.2(8a), RELEASE SOFTWARE
Copyright (c) 1986-1995 by cisco Systems
2500 processor with 14336 Kbytes of main memory
Abort at 0x10200C2 (PC)
>
>h$               Toggle cache state
B [filename]      [TFTP Server IP address | TFTP Server Name]
                  Load and execute system image from ROM or from TFTP server
C [address]       Continue execution [optional address]
D /S M L V        Deposit value V of size S into location L with modifier M
E /S M L          Examine location L with size S with modifier M
G [address]       Begin execution
H                 Help for commands
I                 Initialize
K                 Stack trace
L [filename]      [TFTP Server IP address | TFTP Server Name]
                  Load system image from ROM or from TFTP server, but do not
                  begin execution
O                 Show configuration register option settings
P                 Set the break point
S                 Single step next instruction
T function        Test device (? for help)
Deposit and Examine sizes may be B (byte), L (long) or S (short).
Modifiers may be R (register) or S (byte swap).
Register names are: D0-D7, A0-A6, SS, US, SA, and PC
>
```

The abort message first conveys that the router has aborted and that you successfully halted the router OS. The second indication that you are in the ROM monitor mode is the > prompt. Also in Example C-4, an **h** was entered to display the help listing; this key is the same as the ? key. Most of the ROM monitor is designed for low-level hardware and software debugging, but a couple of commands are worth mentioning:

- **H –** Displays the help messages, as in Example C-4.
- **I** – Initializes the router. It is the same as the **reload** command.
- **$** – Toggles the cache; used for debugging by the TAC.
- **P** – Sets the break point; used for TAC diagnostics.
- **S** – ls a single-step instruction used for TAC diagnostics.

- **T** *function* – Use the ? key behind the **T** command to perform a low-level test of a specific components. This usually performs a detailed hardware memory diagnostic.
- **B** – Allows manual booting from the ROM monitor:
  - **B flash** – Boots the first file in Flash memory.
  - **B** *filename* [**TFTP** *host*] – Boots over the network using TFTP.
  - **B flash** *filename* – Boots the file (filename) from Flash memory.
- **L** – Works the same as the **B** command, but the router will not begin execution of the code.
- **O** – Examines the 16-bit boot register.
- **O/R** *0x0000* – Sets the boot register by using a manual hex setting. For example, O/R 0x2102 will set the register to its default.
- D /*S M L V* – Deposit value *V* of size *S* into location *L* with modifier *M.*
- E /*S M L* – Examines location *L* with size *S* with modifier *M. **E/S*** **2000002** examines the boot register directly from memory.

At this time, you can verify whether you have a router that supports the CONFREG utility or one that supports only basic ROM monitor commands. By looking at the ROM monitor prompt, you can determine this. By keying in the ? command, you can determine whether CONFREG is supported. For example, in Example C-5, notice that the prompt is a >, the greater-than sign. This prompt is a good indication that you might have to use basic ROM monitor commands to change the boot register. One last check is to simply key in the ? command for help, as the example demonstrates.

Example C-5 Another Example of a Successful Break into ROM Monito1; Followed by the ? or Help Command, Showing the Presence of the CONFREG Utility

```
Abort at 0x10200C2 (PC)
>?
$                 Toggle cache state
B [filename]      [TFTP Server IP address | TFTP Server Name]
                  Load and execute system image from ROM or from TFTP server
C [address]       Continue execution [optional address]
D /S M L V        Deposit value V of size S into location L with modifier M
E /S M L          Examine location L with size S with modifier M
G [address]       Begin execution
H                 Help for commands
I                 Initialize
K                 Stack trace
```



```
L [filename]      [TFTP Server IP address | TFTP Server Name]
                  Load system image from ROM or from TFTP server, but do not
                  begin execution
O                 Show configuration register option settings
P                 Set the break point
S                 Single step next instruction
T function        Test device (? for help)

Deposit and Examine sizes may be B (byte), L (long) or S (short).
Modifiers may be R (register) or S (byte swap).
Register names are: D0-D7, A0-A6, SS, US, SR, and PC
>
```

Example C-6 shows the output from the ? command showing the CONFREG utility. Therefore, to configure this router's boot register, you use CONFREG. Notice in Example C-6 the prompt of **rommon**. This is a good indication that CONFREG is supported.

Example C-6 The ? Command Used on a Router That Supports CONFREG

```
*** System received an abort due to Break Key ***
signal= 0x3, code= 0x0, context= 0x6033f2bB
PC = 0x6005eba4, Cause = 0x20, Status Reg = 0x34408302
rommon 1 >
rommon 1 > ?
alias             set and display aliases command
boot              boot up an external process
break             set/show/clear the breakpoint
confreg           configuration register utility
cont              continue executing a downloaded image
context           display the context of a loaded image
cookie            display contents of cookie PROM in hex
dev               list the device table
dir               list files in file system
dis               disassemble instruction stream
dnld              serial download a program module
frame             print out a selected stack frame
help              monitor builtin command help
history           monitor command history
meminfo           main memory information
repeat            repeat a monitor command
reset             system reset
set               display the monitor variables
stack             produce a stack trace
sync              write monitor environment to NVRAM
```

```
sysret            print out info from last system return
unalias           unset an alias
unset             unset a monitor variable
rommon 2 >
```

At times, reading the English wording of CONFREG can actually be harder to understand than just manipulating the bits in the register. To help understand which bits the questions in CONFREG correspond to, consult Table C-5.

Table C-5: CONFREG to BIT Comparison

CONFREG Text	Bit(s) Set	Default Setting
`enable "diagnostic mode"? y/n [n]:`	15	Off
`enable "use net in IP bcast address"? y/n [n]:`	14	Off
`disable "load rom after netboot fails"? y/n [n]:`	13	On
`enable "use all zero broadcast"? y/n [n]:`	10	Off
`enable "break/abort has effect"? y/n [n]:`	8	Off
`enable "ignore system config info"? y/n [n]:`	6	Off
`change console baud rate? y/n [n]:`	11&12	Off and Off
`change the boot characteristics? y/n [n]:`	0-3	0x2

Password Recovery: Routers

When you have a solid understanding of how the boot register works, password recovery becomes straightforward. For all the router platforms, the procedure involves simply changing bit 6, which ignores the startup-config in NVRAM, and then reloading the router. When the router reboots, it will no longer have a running-config. The configuration is still stored in NVRAM and can be viewed by performing the **show startup-config** command from Privileged mode. Because there is no running-config, there will be no enable password. Therefore, you can enter Enable mode and copy the startup-config to the running-config, with the **copy startup-config running-config** command. At this time, remember to change the register back, set the enable password, bring up the interfaces (which will be down), and save the new configuration. This entire process is outlined in the step list that follows.

As mentioned previously, the router will always accept a break signal if sent during the first 60 seconds of initialization, regardless of whether bit 8 is set. With this in mind, the following procedure will recover most routers:

Step 1 Attach a PC or PDA with terminal-emulation software to the router's console port through a Cisco rolled cable.

Step 2 Power-cycle the router.

Step 3 Issue a break signal by pressing the **Break** key, or by executing one of the other ways mentioned, within 60 seconds of initialization.

Step 4 Determine what type of ROM monitor you have. Is CONREG supported?

- If Basic ROM monitor:

- Set bit 6: >**o/R 0x2142**. This will set bit 6. Reload the router with the **Initialize** command.

- If CONFREG is supported:

  Run the CONFREG utility: >**CONFREG**. Answer every question with the default or Enter, until you come to the question: Enable **ignore system config info**. Answer "yes" to this question. This will also set bit 6. Reload the router with the **RESET** command.

Step 5 When the router reloads, it will try to run setup. Abort the setup utility with a **Ctrl-C**.

Step 6 Enter Privileged mode and do a copy startup-config running-config.(e.g., router# **copy startup-config running-config**).

Step 7 Enter the configuration mode, and do the following:

- Set the boot register back to its original configuration.

- All interfaces will be shut down; bring up all interfaces to their normal state.

- Set the enable password to a new value.

- Save the new configuration.

Caution Be careful after you have ignored NVRAM and reloaded the router. The router still has a configuration in NVRAM, and it is easy to overwrite this configuration with a slip of a keystroke. This is particularly easy for people of the "old school" – a simple **wr** instead of **wr t** will ruin the config stored in NVRAM.

Note Make a backup copy of the current router configuration when modifying the registers or performing any work that could put the router configuration in jeopardy. Taking the small amount of time that it requires to perform this could be priceless if disaster strikes.

Password Recovery: Switches

Password recovery with switches is a little easier than with routers. During the first 30 seconds of initialization, the password and enable password is simply the Enter key. To recover a password on a Catalyst switch, follow this procedure:

Step 1 Power-cycle the switch.

Step 2 As soon as the switch loads, enter Enable mode. This is done by quickly typing in **enable [Enter].** The switch will prompt you for a password. During the first 30 seconds, the password is the Enter key. Therefore, simply press the **Enter** key. In Enable mode, set a new password with the **set password** command. When you are prompted for the old password, use the **Enter** key again.

Step 3 In Enable mode, set a new enable password with the **set enablepass** command. When setting the enable password, you will be prompted for the old password; again, this is simply the **Enter** key.

Upgrading the Cisco IOS Software

At some time, you will have to upgrade the router's Cisco IOS Software. Upgrading Cisco IOS Software is a task that can be trivial if you know what you are doing. The Cisco IOS image is stored on Flash memory, either in SIMMs or in credit-card modules. There are four items to account for before upgrading your router's Cisco IOS Software:

- The router Cisco IOS release-must be Release 9.0 or later. (If this rule applies to you, it might also be time to upgrade to IP version 4.)
- The amount of free space available on Flash.
- The size of the new image, including its DRAM requirements.
- A reachable IP address or name of the server to load the image from.

To locate the amount of Flash space available on SIMMs, simply execute the **show flash** command. To view the contents on a credit-card module, enter **dir** [*device*]- **dir slot()**: and/or **dir slotl**:, depending on which slot has the credit-card Flash. Then use the common Flash commands and their PCMCIA equivalents:

- **show flash** – Displays flash on *SIMMs,* as in Example 1-8.
- **dir [/all | /deleted |** ***/1ong*****]**[*device*][*filename*].
  - /**all** – Lists deleted, undeleted, and files with errors
  - /**deleted**-Lists deleted files only
  - /**long**-Lists files in a long, detailed format
  - *device* – Lists files on a specific Flash device: FLASH:, BOOTTFLASH:, SLOT0:, SLOT1: ,
  - *filename* – Names a specific Flash file to list
- **cd** – Changes from one Flash device to another.
- **copy** *source-device:filename destination-device:filename* – Copies files from one source to another. If no specific file is listed, you will be prompted later to enter the filename. This is the case when you copy TFTP to Flash.

Appendix D: Testing DLSw+

Overview

While there are various **show** and **debug** commands available to verify the operation of DLSW+ in a lab environment, the only true way to test DLSw+ it is to generate traffic across the circuits. This section will discuss different ways to generate various types of traffic in a DLSW+ test environment.

Generating NetBIOS traffic

To generate NetBIOS traffic, you need a couple of old Windows PCs. Anything running Windows 95 or later will work. Attach one PC to each side of the DLSW+ circuit as shown in the diagram below.

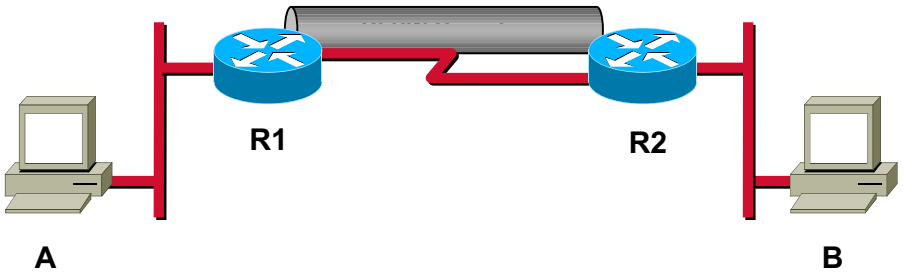

The PCs can be equipped with either Ethernet or Token Ring NICs. It is probably a good idea to have a mixture of both, so that you can test different scenarios. Now, all you need to do is create a share on one PC. If DLSW+ is setup and functioning correctly, you should be able to browse Network Neighborhood on the other PC and see this share. Another quick test is the Find Computer option in Network Neighborhood to locate the other computer by NetBIOS name.

This setup will allow you to actually see NetBIOS traffic across the DLSW+ circuits. This comes in handy, especially when practicing things like NetBIOS name filtering, LSAP filtering, ICANREACH statements, etc.

Generating SNA traffic

Most people do not have spare Mainframes lying around the house, but there are other ways to generate SNA traffic. If you have a couple of spare routers in your lab that meet the memory requirements, you can run the SNASW IOS code on them. This code allows you to emulate a SNA PU and a SNA host. Below is a sample configuration for setting this up.

This sample configuration requires 4 Token Ring routers and can be used to test SNA by simulating a DSPU host and DSPU PU.

Topology:

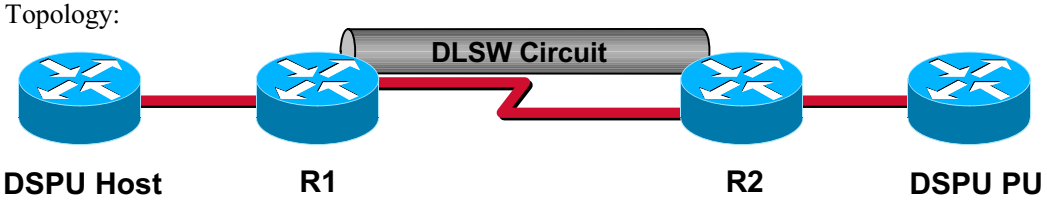

Key points:

- DSPU host - router pretending to be FEP/host side
- DSPU pu - router pretending to be client/end station/PU2.x
- r1 & r2 - routers configured for DLSW+

In this test scenario the DSPU host is configured to accept traffic from a particular client/PU. Once the DSPU PU is defined on the DSPU host, the host sends out a xid that will bring up the DLSW+ circuit.

In a real world environment, a real SNA host sits behind the DSPU host.

Configuration:

(DSPU host)

The DSPU host is configured to recognize the DSPU PU.

```
dspu pu TOK1 rmac 4000.1010.0001 rsap 4 lsap 4


interface TokenRing0
 mac-address 4000.3745.0001
 no ip directed-broadcast
 ring-speed 16
 dspu enable-pu lsap 4
 dspu start TOK1
```


Verify the results:

```
2513#show dspu pu TOK1
dspu pu TOK1 TokenRing0  PU STATUS Active
     RMAC 4000.1010.0001 RSAP 4 LSAP 4
     XID 01010101 RETRIES 4 RETRY_TIMEOUT 30
     WINDOW 7 MAXIFRAME 1472
     ACT BUFFERS 0 SAP ACT BUFFERS 0
     FRAMES RECEIVED 1 FRAMES SENT 1
     LUs USED BY DSPU 0 LUs ACTIVE 0
     LUs USED BY API 0 LUs ACTIVE 0
     LUs ACTIVATED BY HOST BUT NOT USED 0



------------------------------
(R1)

source-bridge ring-group 505
dlsw local-peer peer-id 150.100.20.1
dlsw remote-peer 0 tcp 64.10.1.1
!
interface Loopback0
 ip address 150.100.20.1 255.255.255.0
!
interface TokenRing0
 ip address 150.100.17.1 255.255.255.0
 ring-speed 16
 source-bridge 301 4 505
```

Verify the results:

```
c2504-1#show dlsw peer
Peers:                state     pkts_rx   pkts_tx  type  drops ckts TCP
uptime
 TCP 64.10.1.1       CONNECT        105       102  conf      0    1   0
00:46:43
Total number of connected peers: 1
Total number of connections:     1

c2504-1#show dlsw circuit detail
Index           local addr(lsap)    remote addr(dsap)  state
```

```
uptime
1040187392      4000.3745.0001(04)  4000.1010.0001(04) CONNECTED
00:30:29
        PCEP: 1470A0     UCEP: 5952D0
        Port:To0          peer 64.10.1.1(2065)
        Flow-Control-Tx  CW:20, Permitted:39; Rx CW:20,
Granted:19; Op: Repeat
        Congestion: Low(02), Flow Op: Half: 0/0 Reset 0/0
        RIF = 0630.1F94.12D0
        Bytes:              18/229        Info-frames:           1/1
        XID-frames:          2/4          UInfo-frames:          0/0
Total number of circuits connected: 1

c2504-1#show dlsw reachability
DLSw Local MAC address reachability cache list
Mac Addr         status     Loc.    port                 rif
4000.3745.0001   FOUND      LOCAL   TokenRing0     0630.1F94.12D0

DLSw Remote MAC address reachability cache list
Mac Addr         status     Loc.    peer
4000.1010.0001   FOUND      REMOTE  64.10.1.1(2065)

DLSw Local NetBIOS Name reachability cache list
NetBIOS Name     status     Loc.    port                 rif

DLSw Remote NetBIOS Name reachability cache list
NetBIOS Name     status     Loc.    peer

------------------------------
(r2)

!
source-bridge ring-group 505
dlsw local-peer peer-id 64.10.1.1
dlsw remote-peer 0 tcp 150.100.20.1
!
!
interface TokenRing0/0
 mtu 4464
 ip address 64.10.1.1 255.255.0.0
 ipx network 123
```



```
 ring-speed 16
 source-bridge 80 2 505
```

Verify the results:

```
3620#show dlsw peer
Peers:                state     pkts_rx   pkts_tx  type  drops ckts TCP
uptime
 TCP 150.100.20.1    CONNECT         98       101  conf      0    1   0
00:44:54
Total number of connected peers: 1
Total number of connections:     1

3620#show dlsw circuit detail
Index           local addr(lsap)    remote addr(dsap)  state
uptime
1056964608      4000.1010.0001(04)  4000.3745.0001(04) CONNECTED
00:29:31
        PCEP: 62CB6F24   UCEP: 62CF5B4C
        Port:To0/0          peer 150.100.20.1(2065)
        Flow-Control-Tx  CW:20, Permitted:19; Rx CW:20,
Granted:39; Op: Repeat
        Congestion: Low(02), Flow Op: Half: 0/0 Reset 0/0
        RIF = 06B0.0502.1F90
        Bytes:             229/18         Info-frames:           1/1
        XID-frames:          4/2          UInfo-frames:          0/0
Total number of circuits connected: 1

3620#show dlsw reachability
DLSw Local MAC address reachability cache list
Mac Addr         status     Loc.    port                 rif
4000.1010.0001   FOUND      LOCAL   TokenRing0/0     06B0.0502.1F90

DLSw Remote MAC address reachability cache list
Mac Addr         status     Loc.    peer
4000.3745.0001   FOUND      REMOTE  150.100.20.1(2065) max-lf(4472)

DLSw Local NetBIOS Name reachability cache list
NetBIOS Name     status     Loc.    port                 rif

DLSw Remote NetBIOS Name reachability cache list
```

```
NetBIOS Name     status     Loc.    peer



------------------------------
(DSPU PU)

Here you define the DSPU host that you will try and contact by MAC address

dspu host TOK2 xid-snd 01010101 rmac 4000.3745.0001 rsap 4 lsap 4
!
interface TokenRing0
 mac-address 4000.1010.0001
 ring-speed 16
 dspu enable-host lsap 4
 dspu start TOK2
!
```

Verify the results:

```
c2504-2#show dspu pu TOK2
dspu host TOK2 TokenRing0  PU STATUS Active
     RMAC 4000.3745.0001 RSAP 4 LSAP 4
     XID 01010101 RETRIES 255 RETRY_TIMEOUT 30
     WINDOW 7 MAXIFRAME 1472
     ACT BUFFERS 0 SAP ACT BUFFERS 0
     FRAMES RECEIVED 1 FRAMES SENT 1
     LUs USED BY DSPU 0 LUs ACTIVE 0
     LUs USED BY API 0 LUs ACTIVE 0
          LUs ACTIVATED BY HOST BUT NOT USED 0
```

Here is another sample configuration for testing DLSW+ using 12.1 SNASW code. This configuration requires 4 routers as well.

```
snasw cpname neta.pc1 - set cpname to neta.pc1 - network a, pc 1
!
snasw port TR0 TokenRing0/0 conntype nohpr - creates a Token Ring port with no
hpr routing
!
snasw port E0 Ethernet0/0 conntype nohpr - or creates an Ethernet port with no
hpr routing
!
snasw port S0 Serial0/0 conntype nohpr - or creates a Serial port with no hpr
!
```

```
snasw link TR0PC2 port TR0 rmac 4000.2222.2222 - create a link, Token Ring-to-
Token Ring: PC2 is the other router acting as an SNA node
!
snasw link E0PC2 port E0 rmac 4000.0222.2222 - create a link for Ethernet-to-
Ethernet
!
snasw link EOTR0PC2 port E0 rmac 0200.4444.4444 - create a link for Ethernet-to-
Token Ring
```

Note If an Ethernet device communicates with a Token Ring device, you must convert the destination non-canonical Token Ring address (4000.2222.2222) to canonical Ethernet format (0200.4444.4444)

On the other router - PC2 defines MAC addresses under the appropriate interface type:

```
interface Ethernet 0/0
   mac-address 4000.0222.2222

interface TokenRing0/0
    mac-address 4000.2222.2222
!
snasw cpname NETA.PC2 - define cpname
!
snasw port TR0 Tokenring0/0 conntype nohpr - again define the appropriate ports
!
snasw port E0 Ethernet 0/0 conntype nohpr - again define the appropriate ports
```

Note SNASW links are not required on the receiving routers.

Other commands include:

- **show snasw port**
- **show snasw link**
- **show snasw session**
- **snasw stop**

To start the link creation use the command **snasw start** first on the receiving router and second on the sending router.

E

Appendix E: IOS Services

Overview

This document describes how to use the Cisco HTTP Server

Using the Cisco HTTP Server

The Cisco IOS software includes a HTTP Server that allows you to access the router via a web browser. From this web browser you can issue IOS commands to manage the router.

Cisco HTTP Server Task List

Use of the Cisco HTTP Server is optional. The Cisco HTTP Server is automatically enabled when you use ClickStart to configure a Cisco 1003, Cisco 1004, or Cisco 1005 router. You must enable the Cisco HTTP Server on all other Cisco IOS routers. Once enabled, you will be able to issue Cisco IOS commands on the router using a web browser.

To use and customize the Cisco HTTP Server, you will need to complete the following tasks. The first task is required; the remaining tasks are optional.

Enabling the Cisco HTTP Server

To enable a Cisco router to be configured from a web browser, use the following command in global configuration mode:

Table 1: ip http server Command

Command	Description
`ip http server`	Enables the HTTP Server on a Cisco router. This allows the router to be remotely configured using a web browser.

Now that the Cisco HTTP Server is enabled, you can perform any of the optional tasks or proceed to configure the router using a web browser.

Changing the Cisco HTTP Server's Port Number

By default, the Cisco HTTP Server runs on port 80. To assign the Cisco HTTP Server to a different port, use the following command in global configuration mode:

Table 2: ip http port number Command

Command	Description
`ip http port` ***`number`***	Assigns the port number used to connect to a Cisco router via a web browser.

Controlling HTTP Access to a Cisco Router

To control which hosts can access a router via HTTP, use the following command in global configuration mode:

Table 3: ip http access-class Command

Command	Description
`ip http access-class {access-list-number |name}`	Restricts HTTP access to a Cisco router.

Specifying the User Authentication Method

To specify how HTTP users are authenticated, use the following command in global configuration mode:

Table 4: ip http authentication Command

Command	Description
`ip http authentication {aaa | enable | local | tacacs}`	Specifies the authentication method used to connect to a router via HTTP.

Using a Web Browser to Issue IOS Commands

You can issue most Cisco IOS commands using a web browser. From the router's home page, you click on a hypertext link titled "Monitor the Router." This link takes you to a web page that has a "Command" field. You can type commands in this field as if you were entering commands at the console. This page also displays a list of commands. You can also execute commands by clicking on the hyperlink for the appropriate command.

Accessing Your Router's Home Page

The Cisco IOS software allows users with a privilege level of 15 to access a predefined home page for the router. If you have been assigned a privilege level other than 15, the web page will display and accept only those commands that have been defined for your privilege level.

To access the home page of a router running the Cisco HTTP Server with a default privilege level of 15, perform the following steps:

Step 1 Enter the following command in the URL field of your web browser and press **Return**: http:// router-name/

- For example, to access a Cisco router named router1 at a privilege level of 15, type http://router1/
- Depending on the authentication method that is configured, the web browser will then prompt you for authentication credentials, such as a username and password or enable password.

Step 2 Enter the authentication credentials.

- The browser will display the home page of the router. The router's home page looks something like the Cisco 3640 home page shown below.

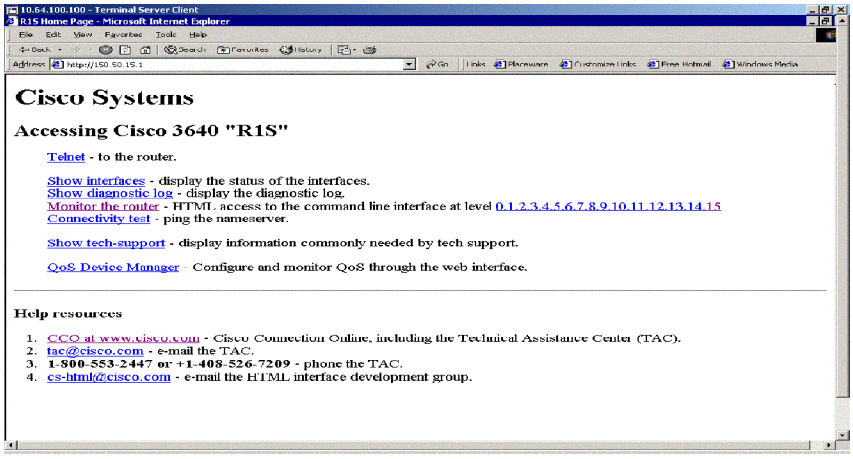

To access a router via HTTP using a privilege level other than 15, perform the following steps:

Step 1 Enter the following command in the URL field of your web browser and press **Return**: http:// router-name/level/level/mode/command.

- For example, to request a user privilege level of 12 on a Cisco router named router1, type http://router1/
- Depending on the authentication method that is configured, the web browser will then prompt you for authentication credentials, such as a username and password or enable password.

Step 2 Enter the authentication credentials.

— The web browser will display a web page specific to your user privilege level, mode, and the command you have requested.

The table below lists the URL arguments that you can use when requesting a web page.

Table 5: URL Arguments

Web Browser URL Arguments	Description
`router-name`	Host name of the router being configured via HTTP.
`level`	The privilege level you are requesting.
`mode`	The mode the command will be executed in, such as exec, configure, interface, etc.
`command`	(Optional) The command you want to execute. If you specify a command, your browser will display a web page showing the results of the requested command. If you do not specify a command in the URL, your browser will display a web page listing hyperlinks of all the commands available for your privilege level.

Issuing Commands via a Web Browser

To issue commands using a web browser, click **Monitor the router** in the first list of hyperlinks on the router's home page. This displays the web page shown below.

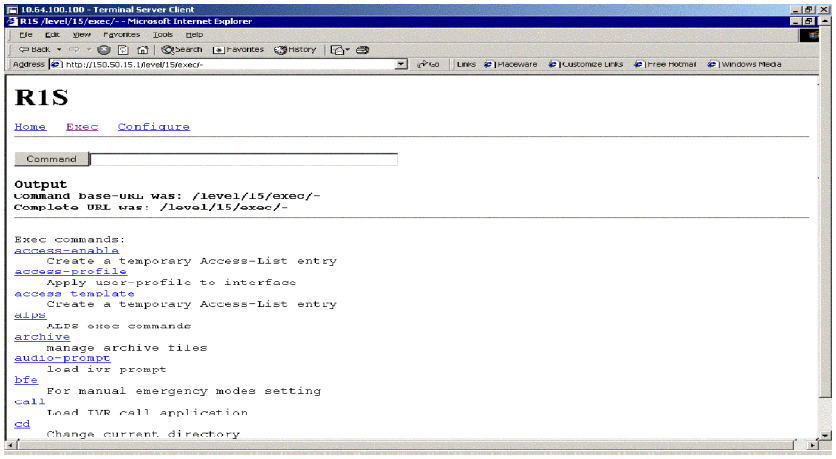

Entering Commands Using Hyperlinks

To enter a command using hyperlinks, scroll through the commands listed at the bottom of the screen and click the one you want to execute. If the link is a complete command, it is executed. If the command has more parameters, another list of command hyperlinks is displayed. Scroll through this second list and click the parameter you want to execute.

If the command is a request for information, like a **show** command, the information is displayed in the web browser.

If the command requires a variable, a form in which you can enter the variable is displayed.

Entering Commands Using the Command Field

Entering the command in the command field is just like entering it at the console. Enter the command using the syntax documented in the Cisco IOS command reference. If you are uncertain of the options available for a particular command, type a question mark (**?**).

For example, entering **show ?** in the command field displays the parameters for the **show** command. The web browser interface displays the parameters as hyperlinks. To select a parameter, you can either click on one of the links, or you can enter the parameter in the command field.

Entering Commands Using the URL

You can also issue a command using the URL field in the web browser.

For example, to execute a **show configuration** command on a router named router2, you would enter the following in the URL window:

http://router2/exec/show/configuration

The web browser then displays the configuration for the router. To save keystrokes, you can actually modify the URL in the URL field for subsequent commands instead of retyping the entire URL.

Caution	There are various security advisories about running the Cisco HTTP Server on a router. You should be fully ware of those before enabling this service on your routers in a production network.

http://www.cisco.com/warp/public/707/ioshttpserverquery-pub.shtml

AutoInstall

The Cisco IOS software includes two features that simplify or automate the configuration of Cisco devices. AutoInstall allows a network manager to load configuration files onto new Cisco devices automatically. Setup is a Cisco IOS software feature that guides a user through the first-time configuration of a Cisco device.

AutoInstall allows you to connect a new router to the network, turn on the new router, and have it configured automatically from a preexisting configuration file. This process was designed to facilitate the centralized management of router installation.

The AutoInstall process begins any time a Cisco IOS software-based device is turned on and a valid configuration file is not found in nonvolatile random-access memory (NVRAM). A configuration file is typically not available when a router is turned on for the first time, or when the configuration file has been manually deleted from NVRAM.

There are two basic approaches to preparing your network for AutoInstall. One approach is to create a minimal configuration file that provides just enough configuration information to allow you to Telnet to the new router and configure it manually. The other approach is to create a host-specific configuration file for each new router containing all of the necessary configuration information. In each case, the configuration file should be created and stored on a TFTP server on the network prior to connecting the new router.

Before the new router can attempt to download a configuration file, however, it must acquire an IP address. This means that a service must be available on the network to provide an IP address to the new router. Your choice of service will determine which interface port on the new router should be connected to the network.

AutoInstall is supported over Ethernet, Token Ring, and FDDI interfaces for LANs and serial interfaces with High-Level Data Link Control (HDLC) encapsulation or serial interfaces with Frame Relay encapsulation for WANs.

If a LAN interface is used, AutoInstall will attempt to acquire an IP address for the attached interface using Dynamic Host Configuration Protocol (DHCP) requests, Bootstrap Protocol (BOOTP) requests, or Reverse Address Resolution Protocol (RARP) requests.

If a serial interface with HDLC encapsulation is connected, AutoInstall will attempt to acquire an IP address for the attached interface using Serial Line Address Resolution Protocol (SLARP). The following table summarizes this information.

Table 6: Protocols Used for IP Address Acquisition in AutoInstall Interface Type

Protocol Used for IP Address Acquisition	Protocol Used for AutoInstall
Ethernet, Token Ring, or FDDI interface	DHCP, BOOTP, or RARP
Serial interface using HDLC	SLARP
Serial interface using Frame Relay	BOOTP

When the AutoInstall process begins, the new router will send DHCP, BOOTP, and RARP requests out any attached interfaces. AutoInstall will use the first available method for configuration. If all LAN interface requests fail, AutoInstall will attempt to configure an available serial interface using SLARP.

Of Token Ring interfaces, only those that set ring speed with physical jumpers support AutoInstall. AutoInstall does not work with Token Ring interfaces for which the ring speed must be set with software configuration commands. If the ring speed is not set, the interface is set to shutdown mode.

In addition to a TFTP server, and a DHCP, BOOTP, RARP, or SLARP server, you may need to configure other elements in your network to enable AutoInstall, as follows:

- If the new router is not directly connected to the device providing the IP address resolution service, you will need to configure the intermediate router to forward requests. This

documentation refers to this intermediate router as the *staging router*. For serial interfaces, a directly attached router providing a SLARP service is required.

- If you wish to enable the new router to download a host-specific configuration file, you can configure a Domain Name System (DNS) server on the LAN network to provide the new router with its hostname. In this case, an IP address-to-hostname mapping for the new router must be added to the DNS database file prior to beginning AutoInstall. Note that a DNS server is not necessary if you configure a DHCP server to provide a hostname for the new router and to provide the IP address of the TFTP server to the new router.

For example, observe the following diagram:

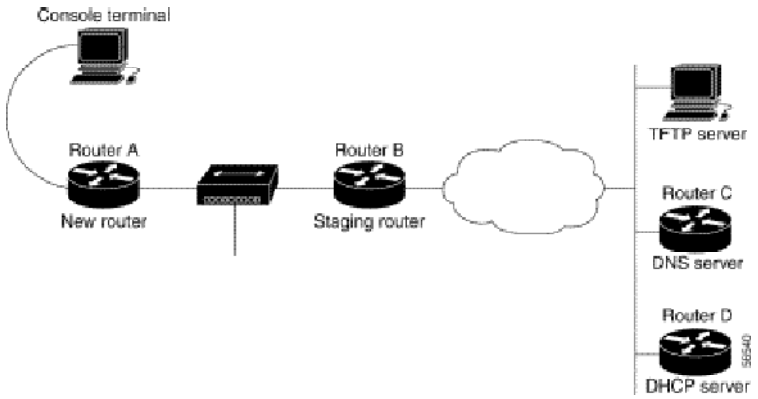

The following steps outline an example of the AutoInstall process that would be used:

1. Router A (the new router) sends a DHCP request out of its attached Ethernet 0 (E0) interface.
2. Router B (the staging router) forwards the request to Router D, which is running a DHCP service.
3. The DHCP server in Router D sends a reply back to Router A. The reply contains a temporary IP address for the E0 interface on Router A and the IP address of the TFTP server.
4. Router A sends a request for a network configuration file to the TFTP server using the address acquired in Step 3.
5. The network configuration file downloaded from the TFTP server does not contain an IP address to hostname mapping for Router A's new IP address, so Router A sends out a DNS request (forwarded by Router B) to acquire its new hostname.
6. The DNS server in Router C resolves the IP address of the new router to the hostname "rtr1" and sends this data to Router A.
7. Router A uses its hostname to send a unicast request to the TFTP server for the host-specific configuration file "rtr1-confg", using the address acquired in Step 3.

8. The "rtr1-confg" file is loaded as the running configuration for Router A. The new configuration contains a permanent IP address assignment, so Router A releases the leased IP address from the DHCP server (using a DHCPRELEASE message).

Note	These steps are simplified to give an impression of the process flow (for example, the DHCP request forwarding in steps 1 and 2 actually consist of multiple discover, offer, and request messages).

Specifying a Staging Router

For those network topologies in which the servers used in the AutoInstall process are not on the same LAN segment as the new router, a staging router is typically needed.

For AutoInstall over serial interfaces, a staging router is always required, because the proper encapsulation must be configured on the interface that will be connected to the new router. Because the address of the staging router cannot be specified to the new router, you should have a direct connection from the new router to the staging router.

The staging router must have an **ip helper** *address* command configured on the appropriate interface for each server to enable unicast requests from the staging router. For example, you may want to configure a helper address for the TFTP server, the DNS server, and the DHCP server on the staging router for AutoInstall over LAN interfaces. For AutoInstall over a Frame Relay encapsulated interface, the staging router will require a helper address for the TFTP server and a Frame Relay map pointing back to the new router.

For AutoInstall over HDLC-encapsulated serial interfaces using SLARP, the interface on the staging router must be configured with an IP address whose host portion has the value 1 or 2. AutoInstall over Frame Relay does not have this address constraint. Subnet masks of any size are supported.

Preparing a Configuration File

Your choice of configuration file will determine many aspects of how you set up the network for AutoInstall. There are three types of configuration files you can make available on a TFTP server:

- A host-specific configuration file, containing a full configuration for each new router ("*hostname*-confg" or "*hostname*.cfg", where *hostname* is a specific router name)
- A default configuration file, containing just enough configuration to allow you to Telnet to the new router and configure it manually ("router-confg", "router.cfg", or "ciscortr.cfg")
- A network configuration file, which allows you to specify IP addresses or hostnames for routers on the network ("network-confg" or "cisconet.cfg")

The syntax of the configuration file-name will depend on the host of the TFTP server. UNIX-based or DOS-based configuration files are saved using the 8.3 naming convention and use ".cfg". Any hostname longer than eight characters is truncated to eight characters. For example, a router with a hostname "australia" will be treated as "australi" and AutoInstall will attempt to download "australi.cfg".

Default network configuration files should have IP address to hostname mappings (using **ip host ip** *address hostname* command line entries).

In general, AutoInstall will attempt to download "-confg" files first, then ".cfg" files. AutoInstall will attempt to download default configuration files in the following order:

- "network-confg"
- "cisconet.cfg"
- "router-confg"
- "router.cfg"
- "ciscortr.cfg"

The request cycle is repeated three times.

AutoInstall over Serial Interfaces

AutoInstall is supported over serial interfaces with HDLC encapsulation or Frame Relay encapsulation. HDLC is the default serial encapsulation. If the AutoInstall process fails over HDLC, it is attempted using Frame Relay encapsulation.

IP Address Acquisition for Serial Interfaces

For AutoInstall over a serial interface, a staging router must be directly connected to the new router using the serial 0 (S0) interface port.

AutoInstall using HDLC

If the new router is connected by an HDLC-encapsulated serial line to the staging router, AutoInstall will send a SLARP request.

In response to a SLARP request, the staging router will send a SLARP reply packet to the new router.

The reply packet contains the IP address and netmask of the staging router's serial interface. If the host portion of the IP address in the SLARP response is 1, AutoInstall will configure the interface of the new router using the value 2 as the host portion of its IP address. If the host portion of the IP address in the SLARP response is 2, AutoInstall will configure the interface of the new router using the value 2 as the host portion of its IP address. The following figure shows an example of this address assignment.

Figure 1: Using SLARP to Assign an IP Address to a New Router

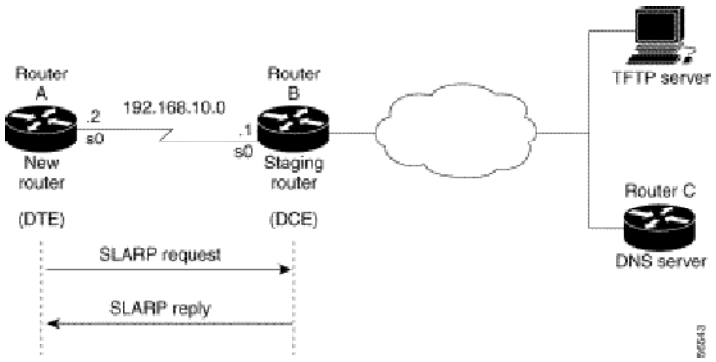

In the figure, the IP address of the Serial 0 interface on the staging router (Router B) is 192.168.10.1. AutoInstall therefore assigns the IP address 192.168.10.2 to the Serial 0 interface of the new router.

Note	If you are using AutoInstall over HDLC, the last 8 bits of host portion of the IP address on the staging router must equal 1 or 2.

AutoInstall Using Frame Relay

If the new router is connected by a Frame Relay-encapsulated serial interface, AutoInstall will send a BOOTP request over the lowest numbered serial or HSSI interface. (The attempt to run AutoInstall over Frame Relay is performed only after attempts are made using SLARP over HDLC, DHCP, and RARP.)

The broadcast BOOTP request sent by the new router will contain the MAC address of the new router's interface. The staging router should be configured to forward the request using a helper address. A DHCP or BOOTP server will then return the IP address assigned to that MAC address. (Note that either a DHCP or BOOTP service can respond to the BOOTP request.)

AutoInstall using Frame Relay can be initiated over only the first serial interface on the new router. Specifically, AutoInstall over Frame Relay can be initiated over Serial 0 (S0), or Serial 1/0 (S1/0). For example, if the new router has serial interfaces S1/0 through S1/3 and S4/0 through S4/3, AutoInstall will be attempted over S1/0 only and cannot be forced to be initiated from S4/0. If AutoInstall over S1/0 fails, a Frame Relay attempt will not be made from any other serial port.

Only a helper address and a Frame Relay map need to be configured on the staging router. No MAC-to-IP address map is needed on the staging router.

Configuration File Downloading for Serial Interfaces

After acquiring an IP address acquired from the RARP, DHCP, or BOOTP server, the new router will attempt to resolve its hostname from a network configuration file or from a DNS service.

The new router will first attempt to resolves its IP address-to-hostname mapping by sending a TFTP broadcast requesting the file "network-confg" or "cisconet.cfg".

The network configuration file is a configuration file generally shared by several routers. In this case, it is used to map the IP address of the new router to the name of the new router. The network configuration file must reside on a reachable TFTP server and must be globally readable. For example, to assign a hostname of "rtr1" to a new router with the address 192.168.10.2, the following line must appear in the network configuration file:

```
ip host rtr1 192.168.10.2
```

If the new router cannot locate and download a "network-confg" or a "cisconet.cfg" file, or if the IP address-to-hostname mapping does not match the newly acquired IP address, the new router sends a DNS broadcast request. If a DNS server is available and has an entry that maps the acquired IP address of the new router to its name, the new router successfully resolves its name.

If DNS does not have an entry that maps the new router's address to its name, the new router cannot resolve its hostname. The new router will then attempt to download a default configuration file ("router-confg", "router.cfg", or "ciscortr.cfg") from the TFTP server. If this attempt also fails, the router will enter Setup mode, or, if using Frame Relay-based AutoInstall, will enter user EXEC mode.

Configuring a Cisco Router as a RARP Server

Use the **ip rarp-server** *ip-address* interface configuration command to enable the RARP service on a Cisco router, where *ip-address* is the IP address of the TFTP server. Use the **arp** *ip-address MAC-address* **arpa** global configuration command to map the MAC address of the new router to a specific IP address.

The following is an example of a static ARP entry in a configuration file for a typical Ethernet host:

```
R5(config)# arp 192.168.7.19 0800.0900.1834 arpa
```

Configuring SNMP

The Simple Network Management Protocol (SNMP) system consists of the following three parts:

- An SNMP manager
- An SNMP agent
- A Management Information Base (MIB)

SNMP is an application-layer protocol that provides a message format for communication between SNMP managers and agents.

The SNMP manager can be part of a Network Management System (NMS) such as CiscoWorks. The agent and MIB reside on the router. To configure SNMP on the router, you define the relationship between the manager and the agent.

The Management Information Base (MIB) is a virtual information storage area for network management information, which consists of collections of managed objects. As explained in RFC 2570: "Collections of related objects are defined in MIB modules. These modules are written in the SNMP MIB module language, which contains elements of OSI's Abstract Syntax Notation One (ASN.1) language. STD 58, RFC 2578, RFC 2579, and RFC 2580 together define the MIB module language, specify the base data types for objects, specify a core set of short-hand specifications for data types called textual conventions, and specify a few administrative assignments of object identifier (OID) values."

The SNMP agent contains MIB variables whose values the SNMP manager can request or change. A manager can get a value from an agent or store a value into that agent. The agent gathers data from the MIB, the repository for information about device parameters and network data. The agent can also respond to a manager's requests to get or set data.

An agent can send unsolicited traps to the manager. Traps are messages alerting the SNMP manager to a condition on the network. Traps can indicate improper user authentication, restarts, link status (up or down), closing of a TCP connection, loss of connection to a neighbor router, or other significant events.

The following figure illustrates the communications relationship between the SNMP manager and agent. A manager can send the agent requests to get and set MIB values. The agent can respond to these requests. Independent of this interaction, the agent can send unsolicited traps to the manager to notify the manager of network conditions.

Figure 2 Communication between an SNMP Agent and Manager

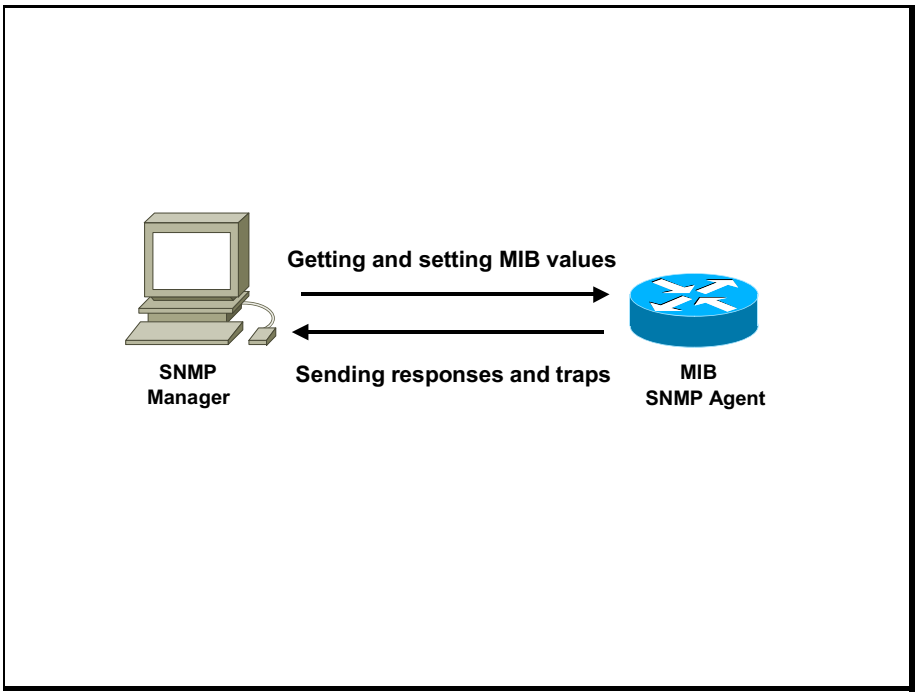

SNMP Notifications

The SNMP Inform Requests feature allows routers to send inform requests to SNMP managers.

Routers can send notifications to SNMP managers when particular events occur. For example, an agent router might send a message to a manager when the agent router experiences an error condition.

SNMP notifications can be sent as traps or inform requests. Traps are unreliable because the receiver does not send any acknowledgment when it receives a trap. The sender cannot determine if the trap was received. However, an SNMP manager that receives an inform request acknowledges the message with an SNMP response PDU. If the manager does not receive an inform request, it does not send a response. If the sender never receives a response, the inform request can be sent again. Thus, informs are more likely to reach their intended destination.

Because they are more reliable, informs consume more resources in the router and in the network. Unlike a trap, which is discarded as soon as it is sent, an inform request must be held in memory until a response is received or the request times out. Also, traps are sent only once, while an inform may be retried several times. The retries increase traffic and contribute to a higher overhead on the network. Thus, traps and inform requests provide a trade-off between reliability and resources. If it is important that the SNMP manager receives every notification, use inform requests. On the other hand, if you are concerned about traffic on your network or memory in the router and you do not need to receive every notification, use traps.

Supported MIBs and RFCs

Cisco's implementation of SNMP supports all MIB II variables (as described in RFC 1213) and defines all traps using the guidelines described in RFC 1215. Cisco provides its own private MIB extensions with every system.

Configuring SNMP

There is no specific command that you use to enable SNMP. The first **snmp-server** command that you enter enables the supported versions of SNMP.

Use an SNMP community string to define the relationship between the SNMP manager and the agent. The community string acts like a password to permit access to the agent on the router. Optionally, you can specify one or more of the following characteristics associated with the string:

- An access list of IP addresses of the SNMP managers that are permitted to use the community string to gain access to the agent.
- A MIB view, which defines the subset of all MIB objects accessible to the given community.
- Read and write or read-only permission for the MIB objects accessible to the community.

To configure a community string, use the following command in global configuration mode:

```
R5(config)# snmp-server community string [view view-name] [ro|rw] [acl-number]
```

Configuring SNMP-Server Hosts

To configure the recipient of an SNMP trap operation, use the following command in global configuration mode:

```
R5(config)# snmp-server host host [traps | informs][version {1 | 2c | 3 [auth |
noauth | priv]} ] community-string  [udp-port port] [notification-type]
```

Establishing the Contact, Location, and Serial Number of the SNMP Agent

You can set the system contact, location, and serial number of the SNMP agent so that these descriptions can be accessed through the configuration file. To do so, use one or more of the following commands in global configuration mode:

```
R5(config)# snmp-server contact text
R5(config)# snmp-server location text
R5(config)# snmp-server chassis-id number
R5(config)#
```


Monitoring SNMP Status

To monitor SNMP status and information, use the following command in EXEC mode:

Table 7: SNMP Commands

Command	Description
R5# show snmp	Monitors SNMP status.
R5# show snmp engineID [local | remote]	Displays information on the local SNMP engine and all remote engines that have been configured on the device.
R5# show snmp groups	Displays information on each SNMP group on the network.
R4(config)# show snmp user	Displays information on each SNMP username in the SNMP users table.

Configuring SNMP Traps

To configure the router to send SNMP traps, use the following commands:

Table 8: SNMP Trap Commands

Command	Description
R#(config)# **snmp-server engineID remote** *remote-ip-addr remote-engineID*	Specify the engineID for the remote host.
R#(config)# **snmp-server user** *username groupname* **remote** *remote-ip-addr* **v3**	Configure an SNMP user to be associated with the above host.
R#(config)# **snmp-server group** [*groupname* {**v1** | **v2c** | **v3** {**auth** | **noauth** | **priv**}}] [**read** *readview*] [**write** *writeview*] [**notify** *notifyview*] [**access** *access-list*]	Configure a group on a remote device.
R#(config)# **snmp-server host** *host-addr* **traps** [version {1 | 2c | 3 [auth | noauth | priv] }] groupname [*notification-type*]	Specify the recipient of the trap message.
R#(config)# **snmp-server enable traps** [*notification-type*] [*notification-option*]	Enable the sending of traps or informs, and specifies the type of notifications to be sent.
R#(config)# **snmp-server manager**	Enable the SNMP manager.

The **snmp-server host** command specifies which hosts will receive traps. The **snmp-server enable traps** command globally enables the trap production mechanism for the specified traps.

In order for a host to receive a trap, an **snmp-server host** command must be configured specifying the intended host, and the trap must be enabled globally through the **snmp-server enable traps** command.

SNMP Examples

The following example enables SNMPv1 and SNMPv2C. The configuration permits any SNMP manager to access all objects with read-only permissions using the community string public. This configuration does not cause the router to send any traps.

```
snmp-server community public
```

The following example permits any SNMP to access all objects with read-only permission using the community string public. The router will also send ISDN traps to the hosts 192.180.1.111 and 192.180.1.33 using SNMPv1 and to the host 192.180.1.27 using SNMPv2C. The community string public is sent with the traps.

```
snmp-server community public
snmp-server enable traps isdn
snmp-server host 192.180.1.27 version 2c public
snmp-server host 192.180.1.111 version 1 public
snmp-server host 192.180.1.33 public
```

The following example allows read-only access for all objects to members of access list 4 that specify the comaccess community string. No other SNMP managers have access to any objects. SNMP Authentication Failure traps are sent by SNMPv2C to the host cisco.com using the community string public.

```
snmp-server community comaccess ro 4
snmp-server enable traps snmp authentication
snmp-server host cisco.com version 2c public
```

The following example sends Entity MIB traps to the host cisco.com. The community string is restricted. The first line enables the router to send Entity MIB traps in addition to any traps previously enabled. The second line specifies the destination of these traps and overwrites any previous snmp-server host commands for the host cisco.com.

```
snmp-server enable traps entity
snmp-server host cisco.com restricted entity
```

The following example sends the SNMP and Cisco environmental monitor enterprise-specific traps to address 172.30.2.160:

```
snmp-server enable traps
snmp-server host 172.30.2.160 public snmp envmon
```

The following example enables the router to send all traps to the host myhost.cisco.com using the community string public:

```
snmp-server enable traps
snmp-server host myhost.cisco.com public
```

The following example will not send traps to any host. The BGP traps are enabled for all hosts, but only the ISDN traps are enabled to be sent to a host.

```
snmp-server enable traps bgp
snmp-server host bob public isdn
```


The following example enables the router to send all inform requests to the host myhost.cisco.com using the community string public:

```
snmp-server enable traps
snmp-server host myhost.cisco.com informs version 2c public
```

F

Appendix F: Answers to Review Questions

Module 2

Module 2: Lesson One Assessment

Q1) What command is used to clear dynamic Frame Relay mappings learned via Inverse ARP?

Answer: clear frame-relay-inarp

Q2) The **frame-relay map** command is used on which of the following interface types?

A) Physical

B) Point-to-multipoint subinterface

C) Point-to-point subinterface

Answer: A. Physical & B. Point-to-multipoint subinterface

A static map links a specified next hop Layer 3 protocol address to a specific DLCI. Partial mesh (hub and spoke) topologies require static maps for spoke-to-spoke communication. Static maps are configured with the **frame-relay map** command.

Q3) The **frame-relay interface-dlci** command is used on which of the following interface types?

A) Physical

B) Point-to-multipoint subinterface

C) Point-to-point subinterface

Answer: C. Point-to-point subinterface

There is no actual remote Layer 3 address-to-DLCI mapping that needs to be configured on a point-to-point subinterface. However, by default, the Frame Relay switch assigns all DLCIs to the physical interface of the Frame Relay DTE. Since each point-to-point subinterface is actually a separate PVC, all you need to do is assign the correct DLCIs to the correct subinterfaces.

Q4) What does the optional **broadcast** keyword on the **frame-relay map** command do?

Answer: The **broadcast** keyword specifies that broadcasts/multicasts (routing updates) should be forwarded across this PVC. You can greatly simplify the configuration for Open Shortest Path First (OSPF) by adding the optional **broadcast** keyword when configuring your FrameRelay map statements.

Q5) Split horizon for IP is disabled on which of the following interface types by default in a Frame Relay topology?

A) Physical

B) Point-to-multipoint subinterface

C) Point-to-point subinterface

Answer: A. Physical and B. Point-to-multipoint subinterface

Due to the NBMA nature of Frame Relay, split horizon for IP is disabled by default on physical and point-to-multipoint subinterfaces.

Module 2: Lesson Two Assessment

Q1) Which of the following indicates a Layer 2 problem?

A) Serial0/0 is up, line protocol is down

B) Serial0/0 is down, line protocol is down

C) Serial0/0 is administratively down, line protocol is down

D) Serial0/0 is up, line protocol is up

Answer: A. Serial0/0 is up, line protocol is down

Q2) Which of the following indicates a Layer 1 problem?

A) Serial0/0 is up, line protocol is down

B) Serial0/0 is down, line protocol is down

C) Serial0/0 is administratively down, line protocol is down

D) Serial0/0 is up, line protocol is up

Answer: B. Serial0/0 is down, line protocol is down

Layer 1 problems are indicated by the following output in the show interfaces command. Serial0/0 is down, line protocol is down

Q2) What command is used to verify Layer 2 connectivity to a directly connected neighbor?

Answer: show cdp neighbors

Q3) Which debug command is used to verify the existence of a Frame Relay map statement when sending pings to a particular next-hop Layer 3 address?

Answer: debug frame packet

Module 2: Lesson Three Assessment

Q1) ATM networks are closely related to which network type?

A) Synchronous

B) Asynchronous

C) Dedicated

D) None of the above

Answer: B. Asynchronous

Q2) Which of the following steps are REQUIRED to configure an ATM connection? (Choose two)

A) Create a PVC

B) Map a protocol address to a PVC

C) Configure the AAL and encapsulation type

D) Configure PVC traffic parameters

Answer: A. Create a PVC & B. Map a protocol address to a PVC

Q2) Configuring ILMI on an ATM connection allows it to discover which type of address?

A) Network layer

B) VPI/VCI

C) DLCI

D) Session layer

Answer: B. VPI/VCI

Q3) Which AAL encapsulation type would you use if you would like to run multiple protocols over a single ATM VC?

A) Aal5snap

B) Aal5mux

C) Aa5encap

D) None of the above

Answer: A. Aal5snap

Module 3

Module 3: Lesson One Assessment

Q1) What is the default encapsulation type on an ISDN BRI interface?

E) PPP

F) HDLC

G) ARPA

H) DDR

Answer: B. HDLC

Q2) Which of the following is an optional component of a dialer profile?

A) Dialer interfaces

B) Dialer pool

C) Physical interfaces

D) Dialer map-class

Answer: D. Dialer map-class

Q2) If access-list 101 is used to specify interesting traffic, which of the following will bring up a DDR link?

```
R4(config)# access-list 101 deny eigrp any any
R4(config)# access-list 101 deny udp any any eq 520
R4(config)# access-list 101 deny tcp any any eq 21
R4(config)# access-list 101 permit ip any any
```

A) RIP

B) FTP

C) EIGRP

D) BGP

Answer: D. BGP

Q3) Which commands should be used on the hub for IP address negotiation? (Pick two)

A) Router(config-if)# ip address negotiated

B) Router(config)# ip local pool default

C) Router(config)# ip address-pool local

D) Router(config-if)# ip unnumbered

Answer: B. Router(config)#ip local pool default 10.0.0.2 10.0.0. 7 & C. Router(config)#ip address-pool local

Q4) Which command is not needed on the physical BRI interface configuration when using dialer profiles?

A) no ip address

B) encapsulation ppp

C) dialer pool-member

D) **dialer-group 2**

Answer: D. **dialer-group 2**

Module 3: Lesson Two Assessment

Q1) Which authentication method sends a clear-text password?

A) CHAP

B) PAP

C) PPP

D) MPPP

Answer: B. PAP

Q2) What authentication mechanism should be used if the destination device supports encrypted hashed messages, but cannot initiate authentication?

A) PAP one-way

B) PAP two-way

C) CHAP one-way

D) CHAP two-way

Answer: C. CHAP one-way

Q2) Which command changes how frequently MPPP calculates the need for additional B channels?

A) ppp timeout multilink link add

B) ppp multilink

C) load-interval

D) dialer load-threshold

Answer: C. **load-interval**

Q3) The “sent-username” feature is used with which two authentication schemes?

A) PAP one-way

B) PAP two-way

C) CHAP one-way

D) CHAP two-way

Answer: A. PAP one-way and B. PAP two-way

Q4) What **CHAP** command should be used on a hub router that requires a different hostname be sent to remote sites?

A) ppp chap altname

B) ppp authentication chap no username

C) ppp chap hostname

D) ppp chap sent-username

Answer: C. **ppp chap hostname**

Module 3: Lesson Three Assessment

Q1) Which backup configuration method uses a static route configured with a higher administrative distance than that of a dynamically learned route to the same location?

A) Backup interface

B) Dialer watch

C) Floating static routes

D) Backup static routes

Answer: C. Floating static routes

Q2) Which backup configuration monitors the status of a route within the routing table?

A) Backup interface

B) Dialer watch

C) Floating static routes

D) None of the above

Answer: B. Dialer watch

Q2) Which routing protocols are supported with dialer watch?

A) RIP

B) OSPF

C) EIGRP

D) BGP

Answer: B. OSPF and C.EIGRP

Q3) Which backup mechanism supports Bandwidth-On-Demand (BOD)?

A) Backup interface

B) Dialer watch

C) Floating static routes

D) All of the above

Answer: A. Backup interface

Q4) Which backup mechanism does not require interesting traffic to initiate a DDR call?

A) Backup interface

B) Dialer watch

C) Floating static routes

D) All of the above

Answer: B. Dialer watch

Module 3: Lesson Four Assessment

Q1) Which **show** command is useful to view the ISDN information for layers 1, 2 and 3?

A) show isdn q921

B) show isdn q931

C) show isdn status

D) show isdn active

Answer: C. **show isdn status**

Q2) Which **show** command can be used to show detailed information about calls in progress?

A) show isdn active

B) show isdn q931

C) **show isdn status**

D) None of the above

Answer: A. **show isdn active**

Q2) What **show** command is useful for the verification of DDR setup?

A) show isdn q931

B) show isdn active

C) show interfaces bri

D) **show dialer interface**

Answer: D. **show dialer interface**

Q3) Which ISDN **debug** command displays data link layer information?

A) debug isdn q921

B) debug isdn q931

C) debug dialer

D) None of the above

Answer: A. **debug isdn q921**

Q4) Which ISDN **debug** command is most appropriate for verifying DDR operation?

A) debug isdn q921

B) debug dialer

C) **debug isdn q931**

D) None of the above

Answer: B. **debug dialer**

Module 4

Module 4: Lesson One Assessment

Q1) The management interface on the Catalyst 3550 belongs to which VLAN by default?

A) VLAN 1

B) All VLANs (it is a trunk port)

C) None – you must create a SVI for VLAN 1 first

D) VLAN 1005

Answer: A. VLAN 1

Q2) Which VTP mode should you use if you wish to configure Extended Range VLANs?

A) Server

B) Client

C) Transparent

D) The Catalyst 3550 does not support Extended Range VLANs

Answer: C. Transparent

Q2) When creating VLANs using the vlan database command, when are your changes actually made to the VLAN database and propagated to other switches in the VTP domain?

A) As soon as the VLAN is created

B) Once you give the VLAN a name

C) Once the switch is rebooted

D) When you enter the **exit** command to go back to privileged exec mode

Answer: D. When you enter the **exit** command to go back to privileged exec mode

Q3) What command can be used to obtain a brief summary of all of the VLANs configured on the switch?

Answer: show vlan brief

Module 4: Lesson Two Assessment

Q1) Which of the following are valid switch port types on the Catalyst 3550?

A) Trunk Ports

B) Tunnel Ports

C) VLAN Ports

D) Hybrid Ports

E) Access Ports

Answer: A. Trunk Ports, B. Tunnel, & E. Access Ports

Q2) List the two commands that are required in interface configuration mode to make a switch port an access port.

Answer: switchport mode access & **switchport access vlan** <vlan id>

Q2) Which of the following commands is used to specify the native vlan on an 802.1Q trunk?

A) **switchport dot1q native** <vlan id>

B) **switchport dot1q trunk native** <vlan id>

C) **dot1q trunk native** <vlan id>

D) **switchport trunk native vlan** <vlan id>

Answer: D. **switchport trunk native vlan** <vlan id>

Q3) The Catalyst 3550 supports which of the following tunneling mechanisms?

A) PPTP

B) IPSec

C) 802.1Q Tunneling

D) Layer 2 Protocol Tunneling

Answer: C. 802.1Q Tunneling & D. Layer 2 Protocol Tunneling

Q4) List the command used in interface configuration mode to turn a Layer 2 switch port into a Layer 3 router port.

Answer: no switchport

Q5) Which of the following protocols facilitates the automatic creation of EtherChannels?

A) Dynamic Trunk Protocol (DTP)

B) VLAN Trunking Protocol (VTP)

C) Port Aggregation Protocol (PAgP)

D) None of the above (EtherChannels must be manually created)

Answer: C. Port Aggregation Protocol (PAgP)

Module 4: Lesson Three Assessment

Q1) Which of the following features shut down a PortFast enabled port when a BPDU is received on that port?

A) RootGuard

B) BPDUGuard

C) LoopGuard

D) 802.1X Guard

E) PAgPGuard

Answer: B. BPDUGuard

Q2) _____ extends SPAN by enabling remote monitoring of multiple switches across your network.

A) SwitchProbe

B) RMON

C) RSPAN

D) Extended SPAN

Answer: C. RSPAN

Q2) With _____ you can forward non-IP packets that the switch does not route between VLAN bridge domains and routed ports.

A) 802.1Q Tunneling

B) Layer 2 Protocol Tunneling

C) InterVLAN routing

D) Fallback Bridging

Answer: D. Fallback Bridging

Module 5

Module 5: Lesson One Assessment

Q1) If the router receives a route update from a RIP neighbor and an internal BGP neighbor for the same route, which one is more believable?

Answer: The RIP route

Q2) If RIP has the passive interface command enabled for an interface, will RIP receive RIP routes on that interface? (Assume there is a downstream RIP device.)

A) Yes

B) No

Answer: A. Yes

Q2) What protocol and port number does RIPv2 use for communication with its RIP neighbors?

A) TCP 500

B) UDP 500

C) TCP 88

D) None of the above

Answer: D. None of the above

Module 5: Lesson Two Assessment

Q1) True or False: If EIGRP passive interface is enabled, EIGRP will still receive routes, but it will not advertise any.

A) True

B) False

Answer: A. True

Q2) True or False: EIGRP is not susceptible to split horizon issues.

A) True

B) False

Answer: B. False

Q2) True or False: EIGRP will not establish a relationship with a neighbor with mismatched timers.

A) True

B) False

Answer: B. False

Q3) By default, EIGRP uses the following metrics on which to base its routing decisions.

A) MTU, Bandwidth, Load

B) MTU, Delay

C) MTU, Bandwidth, Load, Reliability, Delay

D) Bandwidth, Delay

Answer: D. Bandwidth, Delay

Module 6

Module 6: Lesson One Assessment

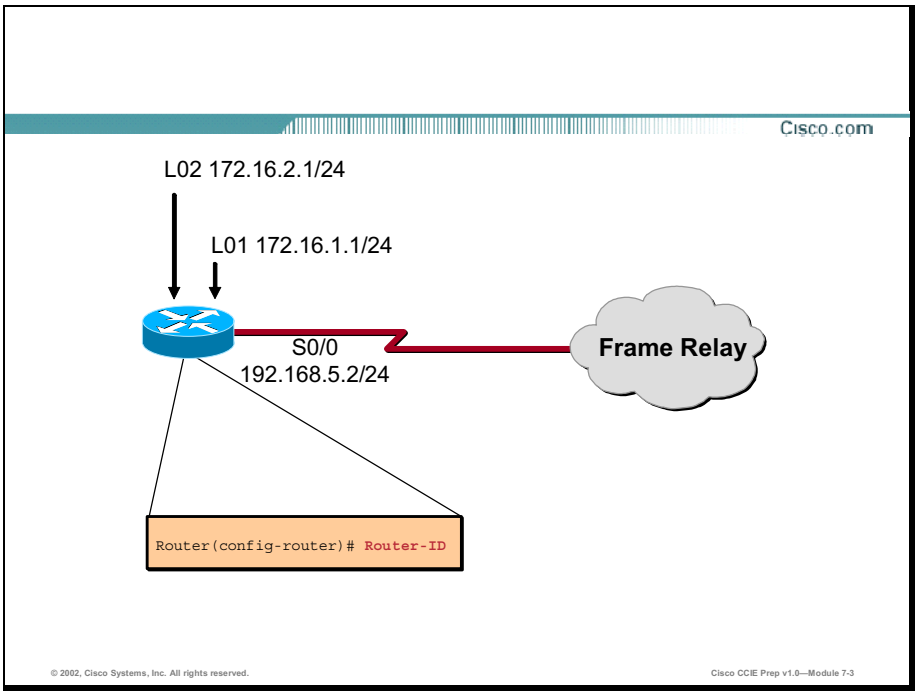

Q1) Based on the configuration above, what will the router ID of this router be?

Answer: The **router-id** command explicitly sets the router ID of the router and overrides all other criteria.

Q2) Which of the following OSPF priority values is used to prevent a router from participating in the DR/BDR election?

A) 0

B) 1

C) 255

D) There is no way to prevent a router from participating in the DR/BDR election

Answer: A. 0

Q2) What command is used to prevent Fast Ethernet and Gigabit Ethernet from both having an OSPF cost of 1?

Answer: auto-cost reference-bandwidth

Q3) Which OSPF network type requires statically defined neighbors and strict control of the DR/BDR election in a hub and spoke NBMA topology?

A) broadcast

B) non-broadcast

C) point-to-point

D) point-to-multipoint

Answer: B. non-broadcast

When using the non-broadcast OSPF network type (default network type for NBMA networks) in a hub and spoke topology, OSPF neighbors must be statically configured using the **neighbor** command. The hub router is also required to become the DR, since it is the only router that has full connectivity to all other routers in the network.

Q4) Which OSPF network types do not require a DR/BDR election?

A) broadcast

B) non-broadcast

C) point-to-point

D) point-to-multipoint

Answer: C. point-to-point and D. point-to-multipoint

Module 6: Lesson Two Assessment

Q1) List the different types of stub areas that Cisco routers support.

Answer: Cisco routers support stub, totally stubby, and not-so-stubby (NSSA) areas.

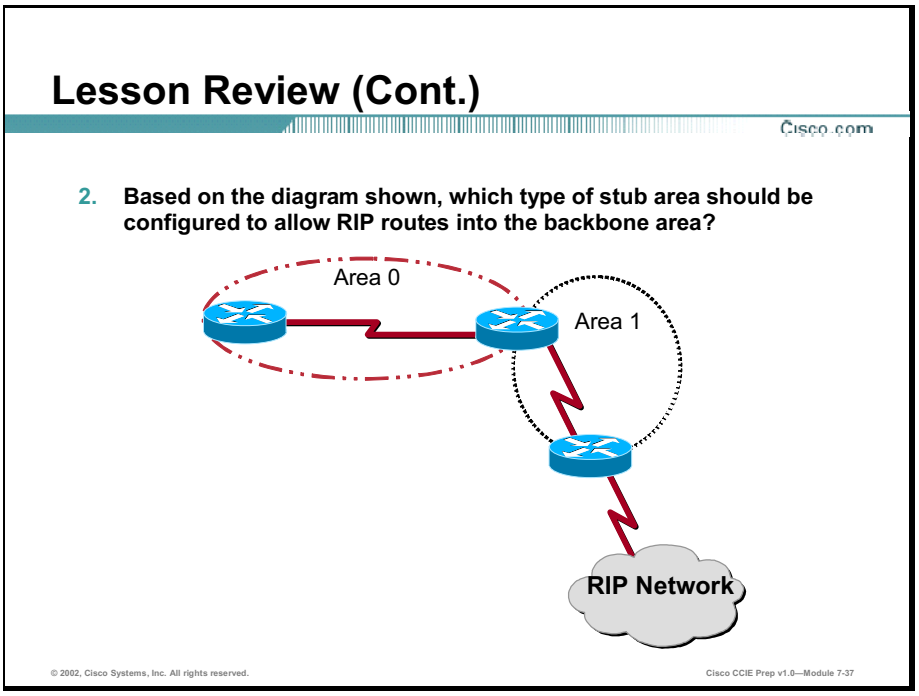

Q2) Based on the diagram above, which type of stub area should be configured to allow RIP routes into the backbone area?

Answer: Not-so-stubby areas (NSSA) are required here because Type 5 LSAs are not allowed in a stub area.

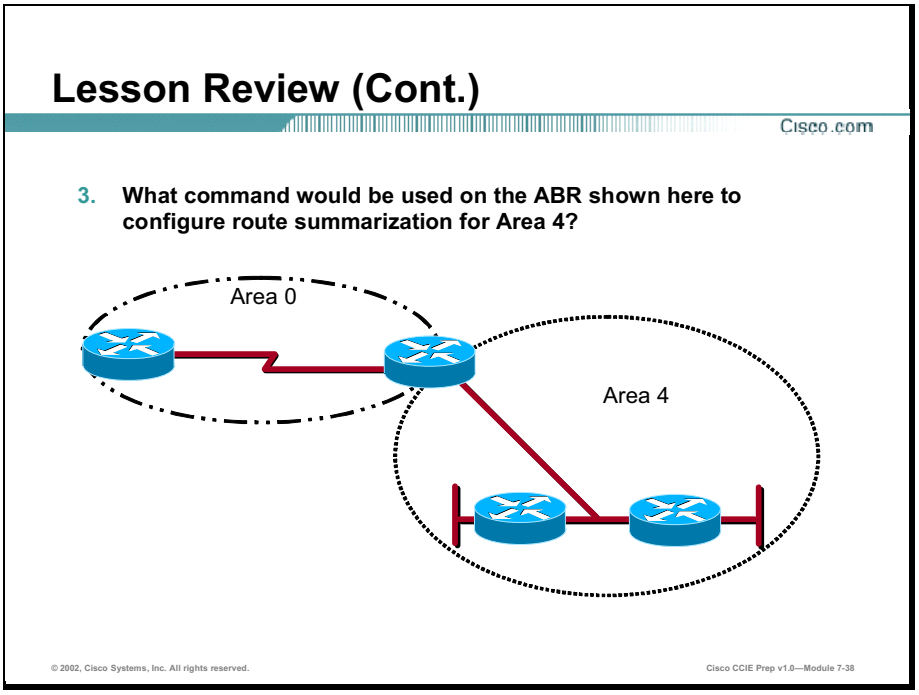

Q2) What command would be used on the ABR shown here to configure route summarization for Area 4?

Answer: Area 4 range

Q3) What command is used to configure external route summarization on an ASBR?

Answer: summary-address. This command instructs the ASBR to summarize external routes before injecting them into the OSPF domain.

Q4) What type of external route increments its cost as it is propagated throughout the OSPF domain?

Answer: External Type 1 routes (E1). External Type 1 routes increment their cost as they pass throughout the OSPF domain.

Module 6: Lesson Three Assessment

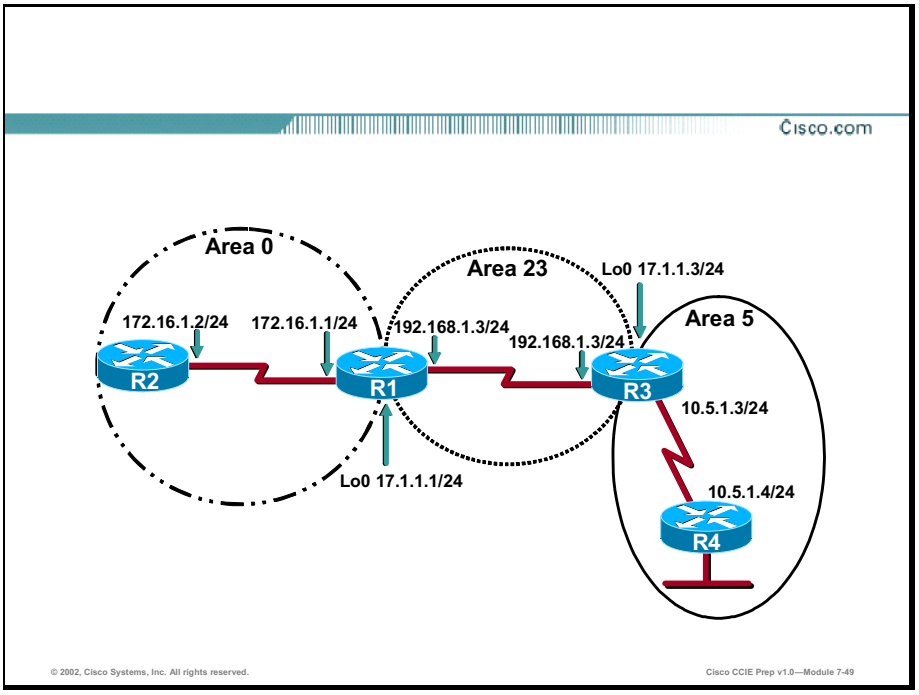

Q1) Based on the diagram shown, what advanced OSPF feature is needed in this network?

Answer: Virtual Link. This diagram represents an area that is not physically connected to Area 0. OSPF requires that all areas be connected to Area 0. Virtual links are used to meet this requirement when it is not physically possible.

Q2) What is the correct command to create a virtual link on R1 in this diagram?

Answer: area 23 virtual-link 17.1.1.1

Q2) In what areas must authentication be configured for R4 in the diagram?

Answer: R4 will need to be configured for Area 0 authentication even though it is not physically attached to Area 0.

Q3) LSAs that have been learned from a neighbor on an OSPF demand circuit are marked as what in the link-state database?

Answer: DNA. The periodic LSA refreshes that take place every 30 minutes in OSPF do not occur over the demand circuit. When the demand circuit is established, a unique option bit (the DC bit) is exchanged between the neighboring routers. If the two routers negotiate the DC bit successfully, they will make a note of it and set a specific bit in the LSA Age field of LSAs they receive from the neighbor on the demand circuit. This specific bit is called the DoNotAge (DNA) bit.

Module 6: Lesson Four Assessment

Q1) What command is used to verify the area in which an interface belongs?

Answer: show ip ospf interface. You can use the **show ip ospf interface** command to verify that OSPF interfaces are running in the correct areas and have the correct OSPF network types defined.

Q2) What command is used to view the OSPF neighbor table?

Answer: show ip ospf neighbor. The show ip ospf neighbor command displays the OSPF neighbor database.

Q2) What command is used to view the router's link-state database?

Answer: show ip ospf database. The **show ip ospf database** command displays the link-state database. The link-state database contains a listing of all the LSAs that a router knows about.

Q3) What command is used to see OSPF neighbor adjacencies, as they are formed in real-time?

Answer: debug ip ospf adj. If an OSPF router is not forming a neighbor adjacency when it should, use the **debug ip ospf adj** command to troubleshoot the adjacency process. This command will display the neighbor adjacency states (DOWN, ATTEMPT, INIT, 2WAY, EXSTART, EXCHANGE, LOADING, and FULL) as they happen in real-time.

Q4) What command is used to verify if an OSPF demand circuit is being brought up due to a change in the link-state topology?

Answer: debug ip ospf monitor. To determine if the link is being brought up due to a change in network topology, use the **debug ip ospf monitor** command. This command shows that LSAs are changing and bringing up the demand circuit.

Module 7

Module 7: Lesson One Assessment

Q1) When your BGP autonomous system ID matches that of your BGP neighbor, what is this considered to be?

A) An EGP relationship

B) External BGP

C) Internal BGP

D) An IGP relationship

Answer: C. Internal BGP

Q2) When running a full mesh iBGP with 10 BGP speakers, how many total peer connections are required?

A) One

B) Four

C) Forty Five

D) Ninety

Answer: C. Forty-Five. The actual formula calculating the number of connections required to maintain a full mesh of point-to-pint link is [n(n-1)/2].

Q2) Using laymen's terms, what does the iBGP synchronization rule state?

A) Any and all routes must be synchronized with the IGP before being placed in the BGP table.

B) Any and all routes must be synchronized with the EGP before being placed in the IP routing table.

C) All BGP peers must have the same (synchronized) BGP table before routes can be placed in the IP routing table.

D) Do not advertise a route if the IGP does not have it in its routing table.

Answer: D. Do not advertise a route if the IGP does not have it in its routing table.

Q3) When creating route reflection for a specific client, on which iBGP peer should the command(s) be placed?

A) The server

B) The client

C) All iBGP peers

D) The hub router in the iBGP

Answer: D. The hub router in the iBGP

Q4) When you have modified an access list used with your BGP neighbor statement, which action would be performed next?

A) Clear the route map

B) Clear the iBGP connections

C) Reload the router

D) Apply the access list to an interface

Answer: B. Clear the iBGP connections

Module 7: Lesson Two Assessment

Q1) True or False. In most situations iBGP neighbors are not directly connected while eBGP neighbors are.

A) True

B) False

Answer: A. True

Q2) Which of the following lessens the full mesh requirement?

A) eBGP multihop

B) confederations

C) communities

D) using loopback interfaces

E) route reflectors

Answer: E. route reflectors

Q2) Which of the following is used to simplify the configuration of a BGP speaker that controls distribution of routing information?

A) eBGP multihop

B) confederations

C) communities

D) using loopback interfaces

Answer: B. confederations

Q3) Which of the following communities is set by default on all destinations?

A) internet

B) no-export

C) no-advertise

D) local-as

Answer: A. internet

Q4) After modifying the community being sent to a neighbor, which of the following commands must also be issued?

A) neighbor <ip-address> send-community

B) neighbor <ip-address> advertise-community

C) clear ip bgp

D) neighbor <ip-address> receive-community

Answer: A. neighbor <ip-address> send-community

Module 7: Lesson Three Assessment

Q1) Which of the following is **NOT** a valid method for advertising a route with Border Gateway Protocol (BGP)?

A) Redistributing static routes

B) Redistributing dynamic routes

C) Redistributing BGP into an Interior Gateway Protocol (IGP) such as Open Shortest Path First (OSPF)

D) Using the **network** command

Answer: C. Redistributing BGP into an Interior Gateway Protocol (IGP) such as Open Shortest Path First (OSPF)

Q2) Which of the following is usually discouraged?

A) Redistributing an IGP into BGP

B) Redistributing BGP into an IGP

C) Redistributing static routes that point to null 0

D) All of the above

Answer: B. Redistributing BGP into an IGP

Q2) When performing redistribution of any kind, which of the following commands is usually required?

A) ip route 0.0.0.0 0.0.0.0 <*ip-address*>

B) default-metric

C) ip classless

D) ip subnet-zero

Answer: B. default-metric

Q3) Which of the following commands would you issue to redistribute EIGRP 10 into BGP Autonomous System (AS) 200?

A) R1(config-router)# redistribute eigrp 10

B) R1(config)# router eigrp 10

C) R1(config-router)# redistribute bgp 200

D) R1(config-router)# default-metric 1000 200 255 1 1500

Answer: A. R1(config-router)# redistribute eigrp 10

Q4) Which command should be issued after modifying your configuration to implement redistribution?

A) default-metric

B) ip route <*ip-address*> <*mask*> null 0

C) clear ip bgp *

D) ip route 0.0.0.0 0.0.0.0 null 0

Answer: B. ip route <*ip-address*> <*mask*> null 0

Module 7: Lesson Four Assessment

Q1) Private Autonomous System (AS) numbers fall into which range?

A) 1-1023

B) 1024-2048

C) 65550-65535

D) 64152 to 65535

Answer: D. 64152 to 65535

Q2) What is the proper term that describes when a Border Gateway Protocol (BGP) prefix is constantly updated and withdrawn from the BGP table?

A) convergence

B) route flapping

C) redistribution

D) dampening

Answer: B. route flapping

Q2) When you wish to perform filtering via Internet Protocol (IP) addresses, which command(s) could you issue?

A) neighbor <*ip-address*> prefix-list

B) neighbor <*ip-address*> distribute-list

C) neighbor <*ip-address*> as-path-list

D) neighbor <*ip-address*> filter-list

Answer: B. neighbor <*ip-address*> distribute-list and D. neighbor <*ip-address*> filter-list

Q3) When you wish to perform filtering via an AS path, which command(s) could you issue?

A) neighbor <*ip-address*> prefix-list

B) neighbor <*ip-address*> distribute-list

C) neighbor <*ip-address*> as-path-list

D) neighbor <*ip-address*> filter-list

Answer: A. neighbor <*ip-address*> prefix-list and C. neighbor <*ip-address*> as-path-list

Q4) Which of the following regular expressions will only allow networks originating from AS 600 to enter a BGP router?

A) ip as-path access-list 1 permit ^600$

B) ip as-path access-list 1 permit $600_

C) ip as-path access-list 1 permit ^600_

D) ip as-path access-list 1 permit _600_

Answer: **A.** ip as-path access-list 1 permit ^600$ and C. ip as-path access-list 1 permit ^600_

Module 7: Lesson Five Assessment

Q1) Which command would you issue to display entries in the Border Gateway Protocol (BGP) routing table?

A) show ip route

B) show ip bgp

C) show ip bgp route

D) show ip bgp summary

Answer: B. show ip bgp

Q2) Which command would you issue to display routes that belong to specified BGP communities?

A) show ip bgp summary

B) show bgp community

C) show communities

D) show ip bgp community

Answer: D. show ip bgp community

Q2) Which command would you issue to display information about BGP peer groups?

A) show ip bgp peer group

B) show bgp peer group

C) show ip bgp peer-group

D) show bgp peer-group

Answer: C. show ip bgp peer-group

Q3) Which debug command would you issue to view output of a BGP speaker making a proper BGP neighbor relationship?

A) show ip bgp

B) debug ip bgp neighbor

C) debug ip bgp

D) debug bgp all

Answer: C. debug ip bgp

Q4) Which debug command would you issue to display BGP dampening?

A) show ip bgp dampening

B) debug dampening

C) debug ip dampening

D) debug ip bgp dampening

Answer: D. debug ip bgp dampening

Module 8

Module 8: Lesson One Assessment

Q1) How can you inject a default route into OSPF?

A) Create a static default route, then redistribute it into OSPF

B) Use the OSPF **default-information originate always** command

C) Create a static default route, and it will automatically find its way into OSPF

D) Create an ABR, and the default route will automatically be injected into the non-backbone area

Answer: A. Create a static default route, and then redistribute it into OSPF, B. Use the OSPF **default-information originate always** command, and D. Create an ABR, and the default route will automatically be injected into the non backbone area

Q2) Router1 is directly connected to the 135.10.2.0/24 subnet. When router1 pings the address of 135.10.3.1, there is no echo reply. What may cause this problem?

A) No default gateway on source or destination

B) Routing problem somewhere between the two devices

C) Router1 has a default route, and the command **no ip classless** is in the configuration

D) There is no remote device that is running IP with the address of 135.10.3.1

Answer: All of the above

Q2) How can you inject a default route into RIP?

A) Use the RIP **default-information originate** command

B) Create a static default route, and it will automatically find its way into RIP

C) RIP does not support advertisement of the default route

D) Use the **ip default-gateway** command

Answer: A. Use the RIP default-information originate command and B. Create a static default route, and that will automatically find its way into RIP

Module 8: Lesson Two Assessment

Q1) Using EIGRP, you notice that your subnets do not show up across the entire network. What can you do to correct this?

A) Manually redistribute from EIGRP into OSPF, modify the summary address, then redistribute back into EIGRP

B) Use the *subnets* option for redistribution

C) Use the *no auto-summarize* option

D) This situation cannot be corrected with today's technology

Answer: C. Use the *no auto-summarize* option

Q2) What are the safe techniques for redistribution of routes, without creating a routing loop?

A) Avoid mutual redistribution

B) Use route maps to only allow specific routes in the redistribution

C) Designate OSPF over ISDN as demand circuits

D) Use snapshot routing

Answer: A. Avoid mutual redistribution and B. Use route maps to only allow specific routes in the redistribution

Q2) On an ASBR you use the **area range** command, but the redistributed RIP routes are not being summarized into OSPF. What would cause this?

A) The **area range** command only works on classful boundaries

B) The *subnets* option should be removed within the redistribution statement

C) The **area range** command only summarized OSPF routes, no redistributed routes

D) OSPF can support VLSM, but routes redistributed from RIP must all use the same mask forever

Answer: C. The **area range** command only summarized OSPF routes, no redistributed routes

Q3) How can you redistribute a 28-bit OSPF route into a 26-bit RIPv1 domain?

A) Create static routes with a 26-bit mask that correlate to the OSPF routes, and redistribute those into RIP

B) Allow OSPF to summarize the 28-bit mask networks into a 26-bit mask using the **area range** command

C) Redistribute the OSPF routes into EIGRP, and allow EIGRP to summarize the routes to a 26-bit route on an interface-by-interface basis

D) Use the **redistribute** command, with the *subnets* option

Answer: A. Create static routes with a 26-bit mask that correlate to the OSPF routes, and redistribute those into RIP

Module 8: Lesson Three Assessment

Q1) On the network, some of the routers receive RIP routes and others do not. What could cause this?

A) A router may be directly connected to all networks

B) Distribute lists may be applied

C) The version of RIP may not be matched either globally or on an interface-by-interface basis

D) Authentication may be set incorrectly on some of the routers

E) The passive interface option may be prohibiting some of the routers from receiving updates

Answer: All of the above

Module 9

Module 9: Lesson One Assessment

Q1) _____ are internetworking devices designed to interconnect LANs to form the appearance of a single larger data link.

A) Bridges

B) Firewalls

C) Routers

D) All of the above

Answer:

Q2) Which of the following technologies allow a protocol to be bridged on one set of interfaces and routed on another set of interfaces?

A) SRB

B) SRT

C) IRB

D) SR/TLB

Answer: C. IRB

Q2) What is the default bridge priority when using the IEEE protocol?

A) 0

B) 32768

C) 128

D) 100

E) 16384

Answer:

Module 9: Lesson Two Assessment

Q1) What are the minimum components of an Ethernet DLSw+ configuration? (Choose all that apply.)

A) Virtual ring-group

B) Pseudo bridge-group

C) Local Peer

D) Remote Peer

Answer: A. Virtual ring-group & C. Local Peer

Q2) Which of the following is NOT a supported DLSw+ encapsulation type?

A) TCP

B) FST

C) Direct

D) STP

Answer: D. STP

Q3) Which of the following permits a casual, any-to-any connection without the burden of configuring the connection in advance?

A) Border peers

B) Peer groups

C) On-demand peers

D) Explorer firewalls

Answer: C. On-demand peers

Q4) What types of caching does DLSw+ use to reduce traffic? (Choose all that apply.)

A) Local

B) Remote

C) Border peer or group

D) On-demand

Answer: A. Local, B. Remote, & C. Border peer or group

Module 9: Lesson Three Assessment

Q1) Which of the following commands allows you to show the status of remote peers?

A) **show dlsw status**

B) **show dlsw peers**

C) **show dlsw peer status**

D) **show status**

Answer: B. **show dlsw peers**

Q2) Which of the following commands will let you determine which SNA or NetBIOS DLSw+ end stations a router has in its cache?

A) **show dlsw peers**

B) **show dlsw cache**

C) **show dlsw reachability**

D) **show dlsw circuits**

Answer: C. **show dlsw reachability**

Q3) Which of the following commands displays the state of all circuits involving a particular MAC address or SAP value?

A) **show dlsw circuits**

B) **show dlsw reachability**

C) **show dlsw peers**

D) **show dlsw cache**

Answer: A. **show dlsw circuits**

Q4) What is the normal state of a DLSw+ circuit when it is successfully connected?

A) MULTIPLE-FRAMES-ESTABLISHED

B) UP

C) CONNECTED

D) PEER-PEER

Answer: C. CONNECTED

Q5) Which of the following commands enables UDP debugging for a remote peer, with an IP address of 1.1.1.6?

A) debug dlsw remote-peer 1.1.1.6 udp

B) debug dlsw peer ip-address 1.1.1.6 udp

C) debug dlsw remote udp 1.1.1.6

D) debug dlsw udp remote ip-address 1.1.1.6

Answer: B. **debug dlsw peer ip-address 1.1.1.6 udp**

Module 10

Module 10: Lesson One Assessment

Q1) After configuring ip multicast routing, which steps must be taken to complete configuration of PIM-DM?

A) router(config-if)# **ip pim dense-mode** (on all interfaces)

B) router(config-if)# **ip pim dense-mode** (only interfaces used in multicasting)

C) router(config-if)# **ip pim dense-mode** and

router(config-if)# **ip igmp** (on all interfaces)

D) No additional configuration

Answer: A. router(config-if)# **ip pim dense-mode** (on all interfaces)

Q2) CGMP requires _____.

A) A router

B) A Cisco Router

C) A CGMP enabled Cisco Router

D) Nothing to run

Answer: C. A CGMP enabled Cisco Router

Q3) True or False: IGMP snooping and CGMP are mutually exclusive.

A) True

B) False

Answer: A. True

Q4) Which command will inform you of the source of multicast packets?

A) **show ip mroute**

B) **show ip mroute source**

C) **show ip multicast source**

D) **debug ip packet**

Answer: B. **show ip mroute source**

Q5) What does (*,G) in the routing table tell about the multicast protocol that is running?

A) DVMRP

B) PIM-DM

C) IGMP

D) PIM-SM

Answer: D. PIM-SM

Module 10: Lesson Two Assessment

Q1) On receiving multiple NTP broadcasts from a variety of NTP servers at different strata levels, which would the router choose for time synchronization?

A) Stratum 15

B) Stratum 0

C) Stratum 255

D) Stratum 1

E) None of the above

Answer: D. Stratum 1

Q2) There are multiple methods to configure NTP on a router. Choose which method is best based on the following information: the router is connected to LAN with 4 NTP servers of different strata levels.

A) router(config)# **ntp server**

B) router(config)# **ntp client**

C) router(config-if)# **ntp broadcast**

D) router(config-if)# **ntp broadcast client**

E) router# **clock set**

F) None of the above, routers are already pre-configured to receive NTP

Answer: C. ntp broadcast

Q3) Diagnose why NTP authentication is failing between RouterA and RouterB, even though they can ping each other and are directly connected.

RouterA	RouterB

`ntp authenticate`
`ntp authentication-key 10 md5 cisco`
`ntp trusted-key 10`
`ntp peer` **`<Router B IP Address>`**
	`ntp authenticate`
	`ntp authentication-key 11 md5 ticktock`
	`ntp trusted-key 11`
	`ntp peer` **`<Router A IP Address>`**

A) The trusted key number is wrong on RouterA

B) The md5 value is wrong on Router A

C) Both md5 and trusted-key are wrong on Router A

D) The Routers are not running service timestamp

Answer: B. The md5 value is wrong on Router A

Q4) A company needs to ensure all routers across the United States provide status in local time, but are synchronized. Which commands should the company use to accomplish this task?

A) Purchase another vendor's routers, for Cisco does not support this function

B) **clock set**

C) **service timestamp**

D) **clock timezone**

E) **clock summer-time**

Answer: D. **clock timezone**

Q4) True or False: NTP is the only way to set time on a router.

F) True

G) False

Answer: B. False

Module 10: Lesson Three Assessment

Q1) True or False: NAT modifies the destination address IP address.

A) True

B) False

Answer: B. False

Q5) If the NAT global inside pool is not seen as a group of IP addresses in the outside interface subnet, what action must be taken?

C) A static routing statement on the 'Natting' router must be made

D) The directly connected router to the outside interface must be routing aware of the IP address(es) in the nat global inside pool IP address(es)

E) The configuration must have overload enabled

F) A PIX Firewall should be used

Answer: B. The directly connected router to the outside interface must be routing aware of the IP address(es) in the nat global inside pool IP address(es)

Q6) Upon using a **show ip nat translation** command, there is only one inside global address being utilized. What has happened?

G) Flow Address Translation is enabled

H) Overload is enabled

I) The nat inside global pool is of one address and overload is enabled

J) The access list for the inside local pool is of one address and overload is enabled

Answer: C. The nat inside global pool is of one address and overload is enabled

Q7) There are 50 inside devices but only 8 valid Internet routable IP addresses. What should the configuration be to allow over 8 inside devices to simultaneously access the Internet?

K) The overload sub-command is required

L) Customer must ask for a larger pool of valid Internet addresses

M) The rotary sub-command is required

N) Use only 8 addresses on the inside devices

O) Define the access list for the inside local pool to only allow 8 devices

Answer: A. The overload sub-command is required

Q8) True or False: NAT supports route maps in 12.0 code.

P) True

Q) False

Answer: A. True

Module 10: Lesson Four Assessment

Q1) If four routers are being configured for the same group and have the default stand by priority, which device becomes the active HSRP router?

A) The router with the highest MAC

B) The router with the lowest MAC

C) The first router configured

D) The router with highest IP address

Answer: C. The first router configured

Q9) To ensure that the router for HSRP Group 44 with the highest priority will be the active router, which command must be added to the configuration?

E) router(config-if)# **standby 44 preempt**

F) router(config-if)# **standby 44 ip**

G) router(config-if)# **standby 44 track**

H) router(config-if)# **standby 44 authenticate**

I) router(config-if)# **standby 44 active**

Answer: A. router(config-if)# **standby 44 preempt**

Q10) To configure HSRP load balancing, between two ISL capable routers, _____.

J) Each router supports the other HSRP group, with each router being the active router for one group

K) Overload is enabled

L) The routing protocol will handle load balancing with maximum paths statement

M) The hosts are divided evenly and point to the appropriate router IP address

Answer: A. Each router supports the other HSRP group, with each router being the active router for one group

Q11) HSRP Tracking provides which features?

N) Performs load balancing

O) Allows hosts to track the HSRP multicast

P) Ensures the active router is available

Q) Reduces the likelihood that a router with an unavailable key interface will remain the active router

Answer: D. Reduces the likelihood that a router with an unavailable key interface will remain the active router

Q12) True or False: HSRP is a routing protocol.

R) True

S) False

Answer: B. False

Module 10: Lesson Five Assessment

Q1) Which DHCP command enables DHCP on the router?

A) router(config)# **ip dhcp pool**

B) router(config)# **network**

C) router(config)# **service ip dhcp**

D) router(config)# **service dhcp**

Answer: A. router(config)# **ip dhcp pool**

Q13) Why should the default router IP address be excluded from the DHCP pool?

E) The DHCP Server does not inadvertently create an address conflict

F) The default gateway is a unique device that receives DHCP from a separate mechanism

G) It provides additional IP addresses for NAT

H) The default-router IP address is used as an option

Answer: A. The DHCP Server does not inadvertently create an address conflict

Q14) DHCP is used to aid in management of _____.

I) Login and passwords

J) Router IP address ranges

K) IP addresses

L) Logins

Answer: C. IP addresses

Module 11

Module 11: Lesson One Assessment

Q1) Passwords can be assigned to which of the following lines?

A) Aux

B) TTY

C) VTY

D) All of the above

Answer: D. Line passwords can be configured on all of the above line types.

Q15) What command would you enter on the VTY lines to allow a user Telneting into the router direct access to privilege mode without entering the enable password?

Answer: **privilege level 15**

Q16) Which type of access list is used to implement Lock and Key?

E) Named

F) Time-based

G) Dynamic

H) Reflexive

Answer: C. Dynamic

Q17) Which are the two TCP Intercept modes supported on an IOS router?

I) Reset

J) Intercept

K) Watch

L) Block

Answer: B. Intercept & C. Watch

Q18) Which of the following are used to encrypt data between two routers?

M) CBAC

N) IPSec

O) TCP Interface

P) GRE

Answer: B. IPSec

Module 11: Lesson Two Assessment

Q1) Which VoIP feature will cause a pre-configured phone number to be dialed when the handset of a phone is lifted?

A) POTS

B) PLAR

C) NUM-EXP

D) PSTN

Answer: B. PLAR

Q19) Which dial-peer configuration command points to all numbers in the 4000 – 4999 range?

B) session target 4xxx

C) session target 4…

D) destination-pattern 4xxx

E) destination-pattern 4…

Answer: D. destination-pattern 4…

Q20) Which VoIP feature will cause a dialed number to be converted into a different number, which is used to actually place a call?

F) POTS

G) PLAR

H) NUM-EXP

E) PSTN

Answer: C. NUM-EXP

Q21) Which command configures a PLAR call to extension 5600?

I) router(config-voiceport)# **connection plar 5600**

J) router(config-dialpeer)# **connection plar 5600**

K) router(config-voiceport)# **plar 5600**

L) router(config-dialpeer)# **plar 5600**

Answer: A. router(config-voiceport)# **connection plar 5600**

Module 11: Lesson Three Assessment

Q1) Which two of the following tools can limit the maximum amount of bandwidth sent over an interface?

A) CQ

B) PQ

C) WRED

D) CAR

E) WFQ

F) FRTS

Answer: D. CAR & F. FRTS

Q22) If Frame Relay Traffic Shaping were configured with a Bc of 5600 bits per time interval and a CIR of 56 kbps, what would be the Tc?

G) 125 ms

H) 100 ms

I) 10 ms

J) 1.25 ms

Answer: B. 100 ms

Tc is calculated by dividing Bc by the CIR. 5600/56000 = 0.1 = 100 ms.

Q23) Which command would configure WRED with a 10 percent discard probability for IP Precedence 3 packets at a queue depth of 256?

K) random-detect precedence 3 120 256 10

L) random-detect precedence 3 64 256 100

M) wred precedence 3 120 256 10

N) wred precedence 3 64 256 100

Answer: A. ramdom-detect precedence 3 120 256 10

The command random-detect precedence 3 120 256 10 specifies a minimum threshold of 120, a maximum threshold of 256, and a mark probability denominator of 10, for IP Precedence 3.

Q24) Which of the following is NOT a valid CAR conform or exceed action?

O) continue

P) drop

Q) retransmit

R) set-prec-continue *new-prec*

S) transmit

T) set-prec-transmit *new-prec*

Answer: C. retransmit

Q25) Which congestion management tool services queues in a round robin fashion?

U) CQ

V) PQ

W) WRED

X) CAR

Y) WFQ

Z) FRTS

Answer: A. CQ

Custom Queuing (CQ) services queues in a round robin fashion and removes a specified number of bytes from each queue.

Answer: A. 1812

Appendix G: Course Glossary

Acronym or Term	Expansion of Acronym
AAL	ATM Adaptation Layer
AARP	AppleTalk Address Resolution Protocol
ABR	Available Bit Rate
ABR	Area Border Router
ACK	Acknowledge
ACL	Access Control List
AD	Administrative Distance
AESA	ATM End System Address
AFI	Authority and Format Identifier
AIP	ATM Interface Processor
ANSI	American National Standards Institute
ARP	Address Resolution Protocol
AS	Autonomous System
AS_SET	Autonomous Systems_Secure Electronic Transaction
ASBR	Autonomous System Boundary Routers
ATM	Asynchronous Transfer Mode
ATM NSAP	Asynchronous Transfer Mode - OSI Network Service Access Point
ATMARP	ATM Address Resolution Protocol
AUTH-ACK	Authenticate-Acknowledge
AUTH-NAK	Authenticate-Not Acknowledged
AUTH-REQ	Authenticate-Request
BACP	Bandwidth Allocation Control Protocol
Bc	Committed Burst
BCMSN	Building Cisco Multilayer Switched Networks
BCRAN	Building Cisco Remote Access Networks
BDR	Backup Designated Router
Be	Excess Burst
BECN	Backward Explicit Congestion Notification
BG	Bridge Group
BGP	Border Gateway Protocol
B-ICI	Broadband Interexchange Carrier Interconnect; a.k.a. B-ISDN Inter-Carrier Interface
BOOTP	Bootstrap Protocol
BPDU	Bridge Protocol Data Unit
BRI	Basic Rate Interface
BSCI	Building Scalable Cisco Internetworks
BVI	Bridge-group Virtual Interface
CAM	Content-Addressable Memory

Acronym or Term	Expansion of Acronym
CAR	Committed Access Rate
CBA	Commonwealth Broadcasting Association
CBAC	Context-based Access Control
CBR	Constant Bit Rate
CBWFQ	Class Based Weighted Fair Queuing
CCIE	Cisco Certified Internetwork Expert
CDP	Cisco Discovery Protocol
CDVT	Cell Delay Variation Tolerance
CGMP	Cisco Group Management Protocol
CGMP	Cisco Group Message Protocol
CHAP	Challenge-Handshake Authentication Protocol
CIDR	Classless Interdomain Routing
CIR	Committed Information Rate
Cisco TAC	Technical Assistance Center
CLI	Command-Line Interface
CLIP	Classical IP
CLNS	Connectionless Network Service
CLP	Congestion Loss Priority
CLR	Cell Loss Ratio
CODEC	COmpression/DECompression
CoS	Class of Service
CPE	Customer Premises Equipment (also known as [DTE])
CPU	Central Processing Unit
CQ	Custom Queuing
CRC	Cyclic Redundancy Check
CS	Carrier Selection
CSIS	Cisco Secure Integrated Software
CS-PDU	Convergence Sublayer Packet Data Unit
CVOICE	Cisco Voice over IP
DCC	Data Country Code
DCE	Data Communication Equipment
DCE	Distributed Computing Environment
DDR	Dial-on-Demand Routing
DE	Discard Eligible
DEC	Digital Equipment Corporation
DHCP	Dynamic Host Configuration Protocol
DHCP/BootP	Dynamic Host Configuration Protocol/Bootstrap Protocol

Acronym or Term	Expansion of Acronym
DISC	"DISConnect"
DLCI	Data-Link Connection Identifiers
DLSw+	Data Link Switching Plus
DM	Disconnect Mode
DM	Dense Mode
DMZ	Demilitarized Zone
DNA	DoNotAge
DNS	Domain Name Server
DNS	Domain Name System
DoS	Denial of Service
DQoS	Deploying Quality of Service in the Enterprise
DR	Designated Router
DSCP	Differentiated Services Code (Control) Point
DSP	Domain Specific Part
DSU	Digital Service Units
DTE	Data Terminal Equipment
DTE/DCE	Data Terminal Equipment/Data Communication Equipment
DTMF	Dual Tone Multi-Frequency
DUAL	Diffusing Update Algorithm
DVMRP	Distance Vector Multicast Routing Protocol
eBGP	external BGP or EBGP
eBGP	exterior Border Gateway Protocol
EGP	Exterior Gateway Protocol
EIGRP	Enhanced Interior Gateway Routing Protocol or Enhanced IGRP
ESI	End System Identifier
FDDI	Fiber Distributed Data Interface
FDX	Full-Duplex
FECN	Forward Explicit Congestion Notification
FIFO	First In First Out
FLSM	Fixed Length Subnet Mask
FRMR	Frame Reject Response
FRTS	Frame Relay Traffic Shaping
FST	Fast-Sequenced Transport
FTP	File Transfer Protocol
FUNI	Frame-based User to Network Interface
FXO	Foreign Exchange Office
FXS	Foreign Exchange Station

Acronym or Term	Expansion of Acronym
Gbps	Gigabits per second
GDA	Global Destination Address
GFC	Generic Flow Control
GGP	Gateway-to-Gateway Protocol
GNS	Get Nearest Server
GRE	Generic Routing Encapsulation
H	Handle
HDLC	High-level Data Link Control
HEC	Header Error Control
HSRP	Hot Standby Routing Protocol
HSSI	High-Speed Serial Interface
HTTP	Hypertext Transport Protocol
IA	Interarea
iBGP	internal BGP
ICD	International Code Designator
ICMP	Internet Control Message Protocol
ICND	Interconnecting Cisco Network Devices
IDI	Initial Domain Identifier
IDP	Initial Domain Part
IEEE	Institute of Electrical and Electronics Engineers
IETF	Internet Engineering Task Force
IGMP	Internet Group Management Protocol
IGP	Internal Gateway Protocol
IGP	Interior Gateway Protocol
IGPM	Internet Group Management Protocol
IGRP	Interior Gateway Routing Protocol
ILMI	Integrated Local Management Interface
InARP	Inverse Address Resolution Protocol
InATMARP	Inverse ATM Address Resolution Protocol
IOS	Input/Output Supervisor
IOS	Internetwork Operating System
IP	Internet Protocol
IPCP	Internet Protocol Control Protocol
IPSec	Internet Protocol Security
IPX	Internetwork Packet Exchange
IPXWAN	IPX Over Various WAN Media
ISDN	Integrated Services Digital Network

Acronym or Term	Expansion of Acronym
IS-IS	Intermediate System to Intermediate System
ISL	Inter-Switch Link
ISO	International Organization for Standardization
ISP	Internet Service Providers
ITU-T	International Telecommunication Union Telecommunication Standardization Sector
LAN	Local Area Network
LANE	LAN Emulation
LCP	Link Control Protocol
LECS	LANE (LAN Emulation) Configuration Server
LLC	Logical Link Control
LLQ	Low Latency Queuing
LMI	Local Management Interface
LS	Link-State
LSA	Link-State Advertisement
LSDB	Link-State Database
LSP	Link-State Packet
MAC	Media Access Control
Mbps	Megabits per second
MBS	Maximum Burst Size
MCR	Minimum Cell Rate
MCTD	Maximum Cell Transfer Delay
MD5	Message Digest Version 5
MED	Multi-Exit Discriminator
MINCIR	Minimum CIR
MOSPF	Multicast OSPF
MPPC	Microsoft Point-to-Point Compression
MPPP	Multilink Point-to-Point Protocol
MSCB	Microsoft Callback Control Protocol
MTU	Maximum Transmission Unit
MUX	Multiplex or Multiplexer
NAS	Network Access Server
NAT	Network Address Translation
NBMA	Non-Broadcast Multi-Access
NCP	Network Control Protocol
NDS	Novell Directory Services
NIC	Network Interface Card
NLM	NetWare Loadable Module
NLSP	NetWare Link Services Protocol

Acronym or Term	Expansion of Acronym
NNI	Network-Network Interface
NSAP	Network Service Access Point
NSSA	Not-So-Stubby Area
NTP	Network Time Protocol
NUM-EXP??	Number Expansion
NVRAM	NonVolatile Random-Access Memory or Non-Volatile RAM
ODI	Open Data-Link Interface
OIL	Outgoing Interface List
OSI	Outgoing Interface List
OSPF	Open Shortest Path First
PAP	Password Authentication Protocol
PBX	Private Branch Exchange
PCR	Peak Cell Rate
PDN	Public Data Network
PDU	Packet Data Unit
PDU	Protocol Data Unit
PIM	Protocol Independent Multicast
PIM-DM	Protocol Independent Multicast – Dense Mode
PIM-SM	Protocol Independent Multicast – Sparse Mode
PING	Packet Internetwork Groper
PIX	Private Internet Exchange
PLAR	Private Line Automatic Ringdown
POTS	Plain Old Telephone Service
ppCDV	Peak-to-peak Cell Delay Variation
PPP	Point-to-Point Protocol
PQ	Priority Queuing
PRI	Primary Rate Interface
PSTN	Public Switched Telephone Network
PT	Payload Type
PVC	Permanent Virtual Connection
Q	Queue count
QoS	Quality of Service
QSAAL	Q.2931 Signaling ATM Adaptation Layer
RED	Random Early Detection
RFC	Requests for Comment
RIF	Routing Information Field
RII	Routing Information Indicator

Acronym or Term	Expansion of Acronym
RIP	Routing Information Protocol
RIP v1	classful routing protocol
RIP v2	classless routing protocol
RIT	Route Information Table
ROM	Read-Only Memory
RP	Rendezvous Point
RPC	Remote Procedure Call
RPF	Reverse Path Forwarding
RSM	Route Switch Module
RSVP	Resource reSerVation Protocol
RTMP	Routing Table Maintenance Protocol
RTO	Retransmission TimeOut
RTP	Reliable Transport Protocol
SAP	Service Advertising Protocol
SAR	Segmentation And Reassembly
SCR	Sustained Cell Rate
SDT	Shared Distribution Trees
SEAL	Simple and Efficient Adaptation Layer
SEL	NSAP Selector
Seq-Num	Sequence Number
SIA	Stuck in Active
SIT	Service Information Table
SLA	Service Level Agreements
SM	Sparse Mode
SMDS	Switched Multimegabit Data Service
SMTP	Simple Mail Transfer Protocol
SNA	System Network Architecture
SNAP	Subnetwork Access Protocol
SNMP	Simple Network Management Protocol
SNMP	Simple Network Management Protocol
SPAN	Switched Port Analyzer
SPF	Shortest Path First
SPT	Shortest-path Trees
SPX	Sequenced Packet Exchange
SR/TLB	Source-route translational bridging
SRB	Source-Route Bridging
SRT	Source Routing Transparent

Acronym or Term	Expansion of Acronym
SRTT	Smooth Round Trip Timer
SSH	Secure Shell
STP	Spanning-Tree Protocol
SVC	Switched Virtual Circuit
SYN	Synchronization
TB	Transparent Bridging
Tc	Committed Time Interval
TCP	Transmission Control Protocol
TCP/IP	Transmission Control Protocol/Internet Protocol
TDM	Time-Division Multiplexing
TEI	Terminal Endpoint Identifier
TFTP	Trivial File Transfer Protocol
ToS	Type of Service
TrBRFs	Token Ring Bridge Relay Functions
TrCRFs	Token Ring Concentrator Relay Functions
TSET	Transform-s+B140et
TTL	Time To Live
TTY	Teletype
UBR	Unspecified Bit Rate
UDP	User Datagram Protocol
UNI	User-to-Network Interface
VBR	Variable Bit Rate Real-Time
VBR-NRT	Variable Bit Rate Non Real-Time
VC	Virtual Channel
VC	Virtual Circuit
VCD	Virtual Circuit Descriptor
VCI	Virtual Channel Identifier
VINES	Virtual Integrated Network Service
VLAN	Virtual LAN or Virtual Local Area Network
VLSM	Variable Length Subnet Mask
VoIP	Voice over IP
VP	Virtual Path
VPI	Virtual Path Identifier
VCI	Virtual Circuit Identifier
VTP	VLAN Trunking Protocol
VTY	Virtual Terminal
WAN	Wide Area Network

Acronym or Term	Expansion of Acronym
WFQ	Weighted Fair Queuing
WINS	Windows Internet Naming Service
WRED	Weighted Random Early Detection
XNS	Xerox Network Systems